


REPORT DOCUMENTATION PAGE					<i>Form Approved</i> OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY) 09/13/2017		2. REPORT TYPE Technical Report			3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE The New Dogs of War: The Future of Weaponized Artificial Intelligence				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Johnson, Brian D.; Vanatta, Natalie; Draudt, Alida; West, Julia R.				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Threatcasting Lab Arizona State University Tempe, AZ 85287					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Army Cyber Institute 2101 New South Post Road West Point, NY 10996					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT This report contemplates the future threats the United States may face from the evolution of artificial intelligence. Focused on weaponized AI—artificial intelligence used as a weapon, intentionally or otherwise—the report examines how next-generation threat actors could use AI and advanced machine learning techniques against the U.S. military, government, industry, and private citizens. The report identified three key threats (surveillance and coercion, the AI weapons factory, and the careless destabilization of national security) as well as actions necessary by academia, non-profits, industry, government, and individuals to mitigate, disrupt, or recover from possible threats.						
15. SUBJECT TERMS threatcasting; artificial intelligence; weaponization; threats; machine learning; weaponized artificial intelligence; future; destabilization						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)	
U	U	U	UU	170		

T H E N E W D O G S O F W A R :

THE FUTURE OF WEAPONIZED ARTIFICIAL INTELLIGENCE



A Threatcasting Report from the
Army Cyber Institute at West Point and
Arizona State University's Threatcasting Lab

T H E N E W D O G S O F W A R :

THE FUTURE OF WEAPONIZED ARTIFICIAL INTELLIGENCE

T H E N E W D O G S O F W A R :

THE FUTURE OF WEAPONIZED ARTIFICIAL INTELLIGENCE

Technical Report from Threatcasting West 2017, a workshop hosted at Arizona State University in conjunction with Army Cyber Institute. Report and workshop produced by Brian David Johnson in conjunction with Alida Draudt, Natalie Vanatta, and Julia Rose West. The photographs showcased in this report were taken on campus at Arizona State University.





“

The fact that these tools in many ways level the playing field, amplifying the power of adversaries that may be (militarily or economically) less capable. Consider the case of Russia and last year's election as an example. And it's a double whammy; it's not clear that the obvious, direct methods of counteracting cyber attacks are effective (you can't un-spill the milk after hacking of a political candidate). This demands new strategies that I don't know are fully in place.

- Sean P. Cornelius, Northeastern University



If we are to outmaneuver our adversaries in cyber space, then we have to be better at imagining the futures they may hope to exploit, and then set the conditions that will ultimately prevent them from harming us. This is not a one-time thing, but a process that gets refined as new information becomes available. We all (government, academia and the private sector) need to ideate and work together to do this right. We can't let this be just an Army thing.

- Fernando Maymi, Soar Technology

Table of Contents

Participants and Lab Team

- 09 Army Cyber Institute Overview
- 09 Arizona State University Threatcasting Lab

About

- 10 The New Dogs of War
- 12 Threatcasting A Brief Overview

Threats

- 17 AI Surveillance and Coercion: The New Dogs Of War
- 18 The AI Weapons Factory
- 19 Careless Destabilization of National Security

Actions

- 23 Academia
- 23 Industry
- 25 Non-profits
- 26 Individuals
- 26 Government

Post Analysis Themes

- 29 The Long Game
- 29 Trust
- 30 Job Loss
- 30 Human-Centric

Current Work and Implications

- 32 AI Surveillance And Coercion: The New Dogs Of War
- 36 The AI Weapons Factory
- 38 Careless Destabilization Of National Security

42 Conclusion

Threat Futures

- 14 The AI Weapons Factory
- 15 A Hole in the Heart
- 21 The Edge of Nothing
- 24 Little Sister Lost
- 27 Trust Fall
- 31 Old Dogs. New Tricks.
- 37 The Automatic Spy
- 41 James and the Giant Truck

07

Appendices

- 44 Threatcasting Explained – Extended Version
- 47 Research Inputs
- 63 Research Synthesis Workbooks (raw)
- 90 Futures Workbook Day One (raw)
- 127 Futures Workbook Day Two (raw)

Participants

Michael G. Bennett	ASU Center for Law, Science and Innovation
Diana M. Bowman	Arizona State University
Jason Brown	US Army
Steve Brown	BaldFuturist.com
Aaron Chan	Office of the Deputy Assistant Secretary of the Army (Research and Technology)
Sean P. Cornelius	Northeastern University
Jeremiah Cottle	United States Secret Service, Office of Strategic Planning & Policy
Don Stanley Dalisay	United States Military Academy
Thomas Dougherty	United States Secret Service, Office of Strategic Planning & Policy
Frank Downs	ISACA
Bridy Godwin	US Army Network Enterprise Technology Command
Christopher Hartley	Army Cyber Institute at West Point
Jason Huff	Lockheed Martin Corporation
Nolan Hedglin	United States Military Academy
Richard Johanning	AECOM
Terence M. Kelley	Army Cyber Institute at West Point
Toby Kohlenberg	
Paul LaDue	US Army Network Enterprise Technology Command
Ryan Lee	United States Air Force Academy
Ian MacLeod	Fractal Industries
Fernando Maymi	Soar Technology
Andrew Maynard	Arizona State University
Sean P. McCafferty	Army Cyber Command
Jack Murray	US Army TRADOC, Army Capabilities Integration Center
Robert Norwood	United States Military Academy
James Pruneski	United States Military Academy
Blake Rhoades	Army Cyber Institute at West Point
Heather Ross	Arizona State University, Global Security Initiative
Gregory Try	United States Secret Service, Office of Strategic Planning & Policy
Austin Yamada	The University of Arizona
Jessica Zhu	United States Military Academy

And many others

ASU Threatcasting Lab Team

Brian David Johnson	Director
Natalie Vanatta	Senior Advisor to the Lab
Alida Draudt	Futurist
Julia Rose West	Futurist



Army Cyber Institute

Unique within the U.S. military, the Army Cyber Institute at West Point is an innovative mix of academic think tank and operational laboratory. ACI's multi-disciplinary team of military, industry, and academic entrepreneurs develop intellectual capital through research and partnerships, enabling the U.S. to outmaneuver its adversaries in cyberspace. Positioned to establish and maintain relationships with the nation's economic center of gravity in New York City, the ACI also directs and synchronizes efforts across the U.S. Military Academy in the cyber domain. The ACI collaborates with the U.S. Army Cyber Command and U.S. Army Cyber Center of Excellence to prevent strategic surprise and ensure the Army's cyber dominance.



Arizona State University Threatcasting lab

The Threatcasting Lab at Arizona State University serves as the premier resource for strategic insight, teaching materials, and exceptional subject matter expertise on Threatcasting, envisioning possible threats ten years in the future. The lab provides a wide range of organizations and institutions actionable models to not only comprehend these possible futures but to a means to identify, track, disrupt, mitigate and recover from them as well. Its reports, programming and materials will bridge gaps, and prompt information exchange and learning across the military, academia, industrial, and governmental communities.

The New Dogs of War

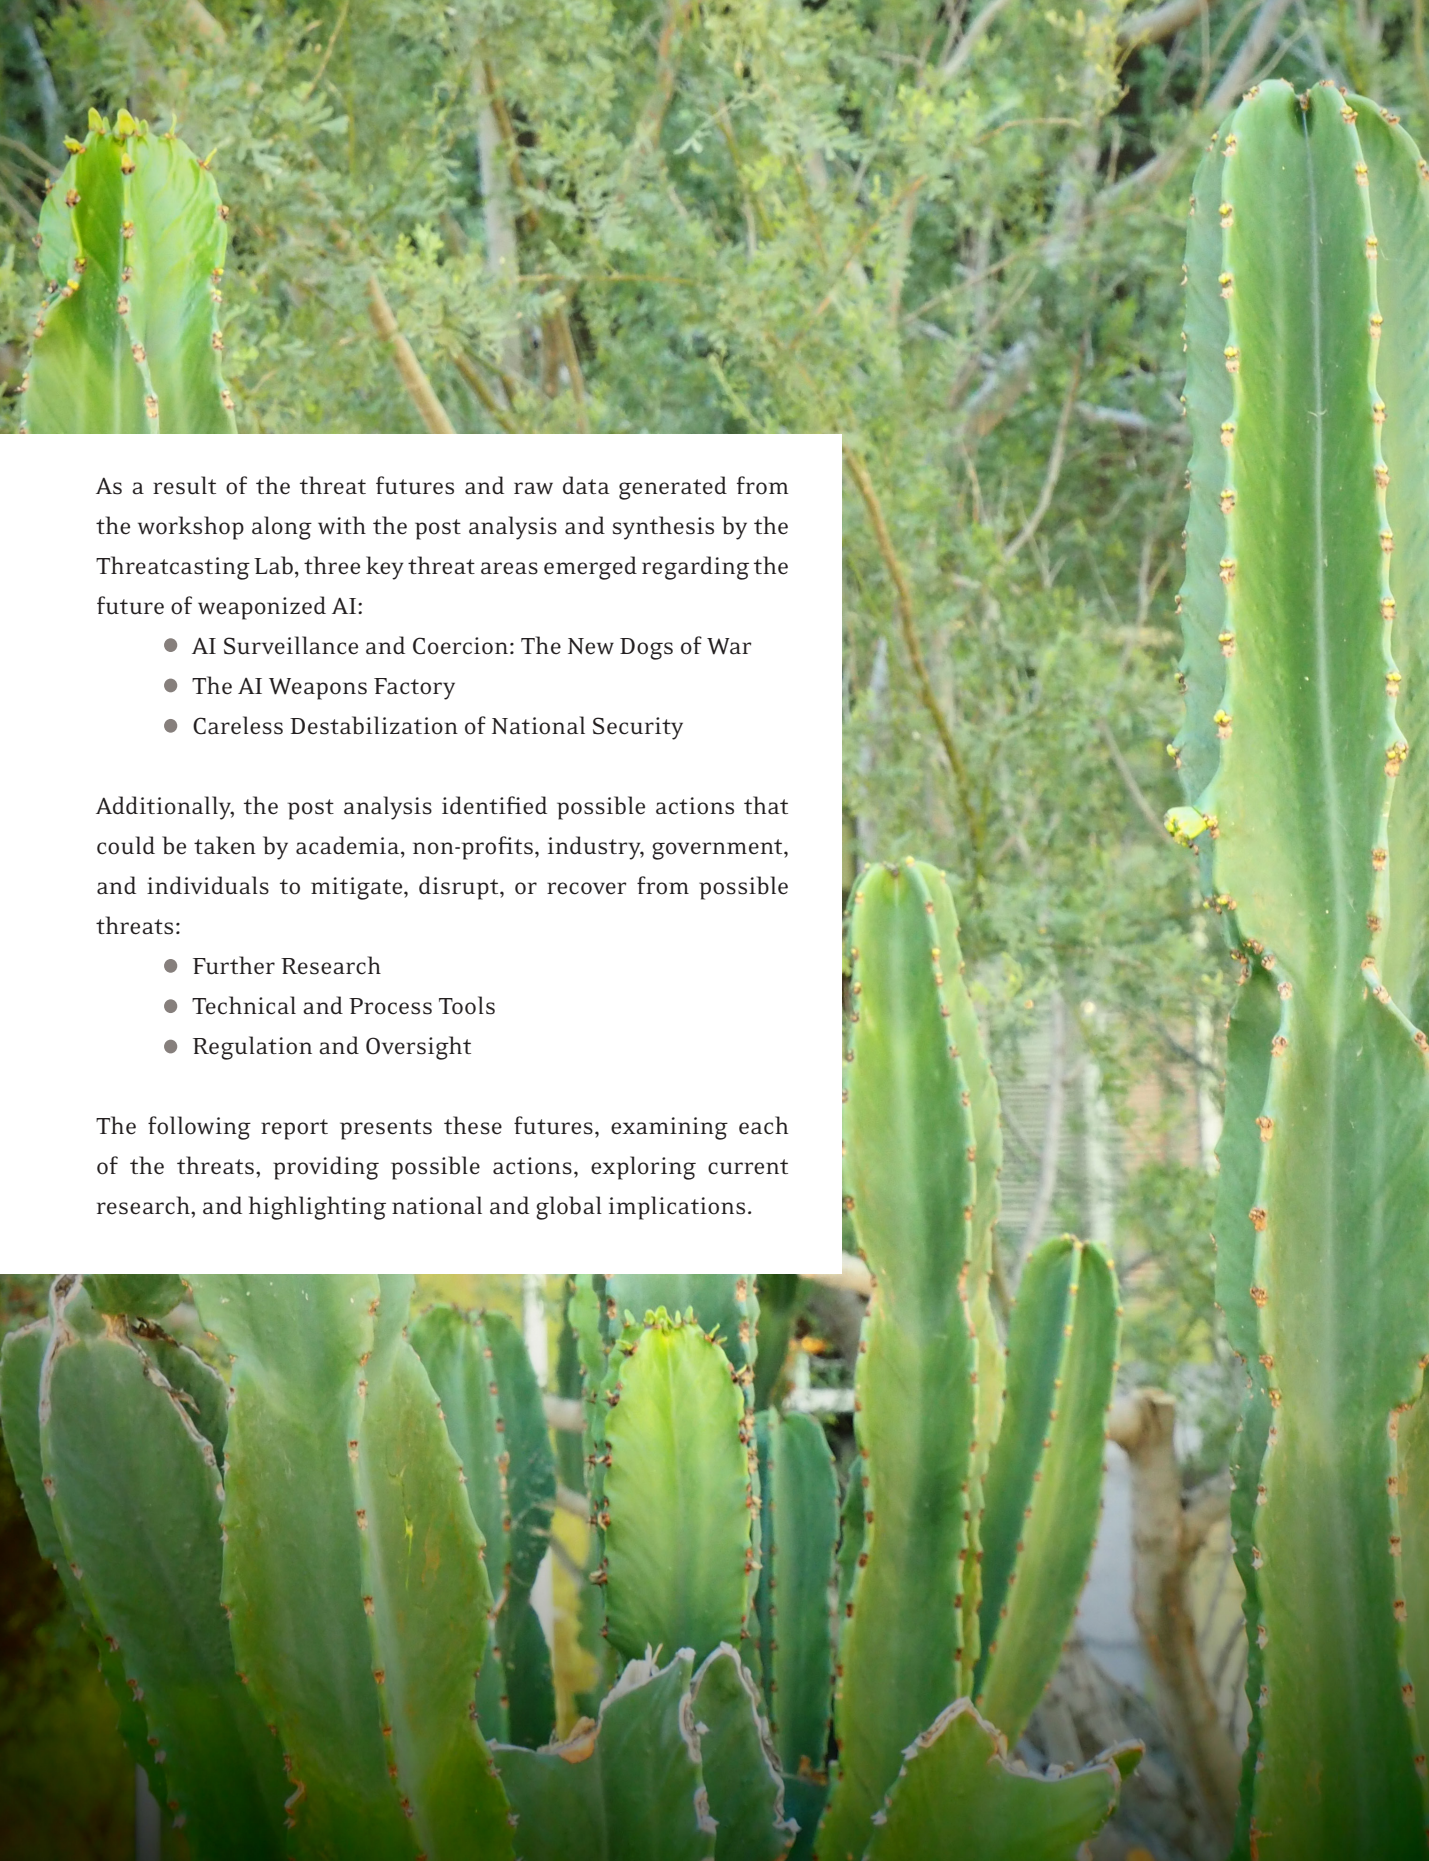
The Future of Weaponized Artificial Intelligence

In May 2017, Arizona State University (ASU) hosted Threatcasting West, a workshop run by the Threatcasting Lab, a joint endeavor between ASU and the Army Cyber Institute. The event brought together individuals from across the military, government, academia and private industry to envision possible threats ten years in the future and what actions can be taken to identify, track, disrupt, mitigate, and recover from possible threats.

A previous Threatcasting East workshop (August 2016, West Point, NY) identified threats resulting from the weaponization of data, including artificial intelligence (AI) and its effect on global supply chains. Threatcasting West 2017 continued this exploration, delving specifically into how next generation threat actors could use AI along with advanced machine learning techniques against the United States military, government, industry and private citizens.

Over the course of two days, 47 participants created 22 unique threat futures, exploring the advancement of AI, the diminishing ability to conduct covert intelligence gathering, the growing complexity of code, and future division of work roles between humans and machines. As a part of the workshop, subject matter experts (SMEs) provided research inputs from which the participants developed their models. This research included the following: How to interrogate and rethink the very nature of AI (Dr. Genevieve Bell), Can we develop AI without losing control over it? (Sam Harris), Cyber considerations for humans and intelligence gathering (Dr. Dave Gioe), How to approach Threatcasting and future modeling from an economic perspective (Paul Thomas), What will be the growth, impact, and future of applying AI to real world industries? (Andre LeBlanc), and A survey of cyber growth and our relationship with machines (MAJ Natalie Vanatta, Ph.D.). Full transcripts of these inputs are located in Appendix, Research Inputs.





As a result of the threat futures and raw data generated from the workshop along with the post analysis and synthesis by the Threatcasting Lab, three key threat areas emerged regarding the future of weaponized AI:

- AI Surveillance and Coercion: The New Dogs of War
- The AI Weapons Factory
- Careless Destabilization of National Security

Additionally, the post analysis identified possible actions that could be taken by academia, non-profits, industry, government, and individuals to mitigate, disrupt, or recover from possible threats:

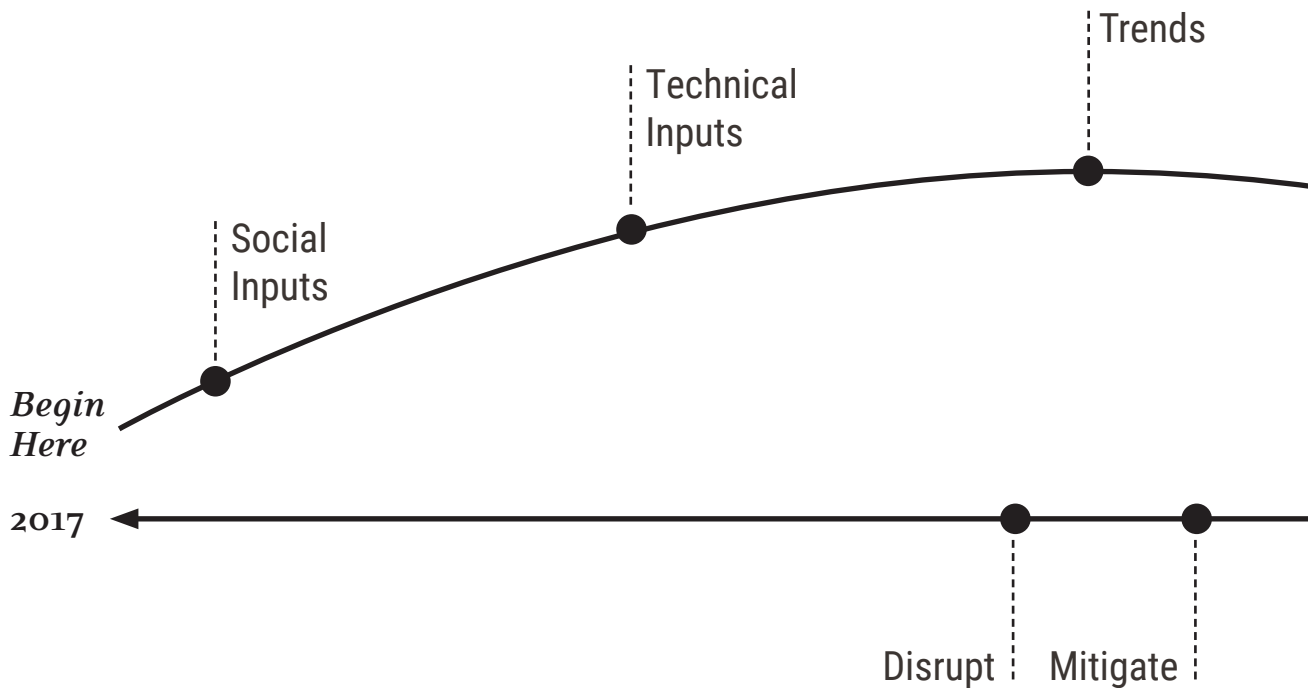
- Further Research
- Technical and Process Tools
- Regulation and Oversight

The following report presents these futures, examining each of the threats, providing possible actions, exploring current research, and highlighting national and global implications.

Threatcasting

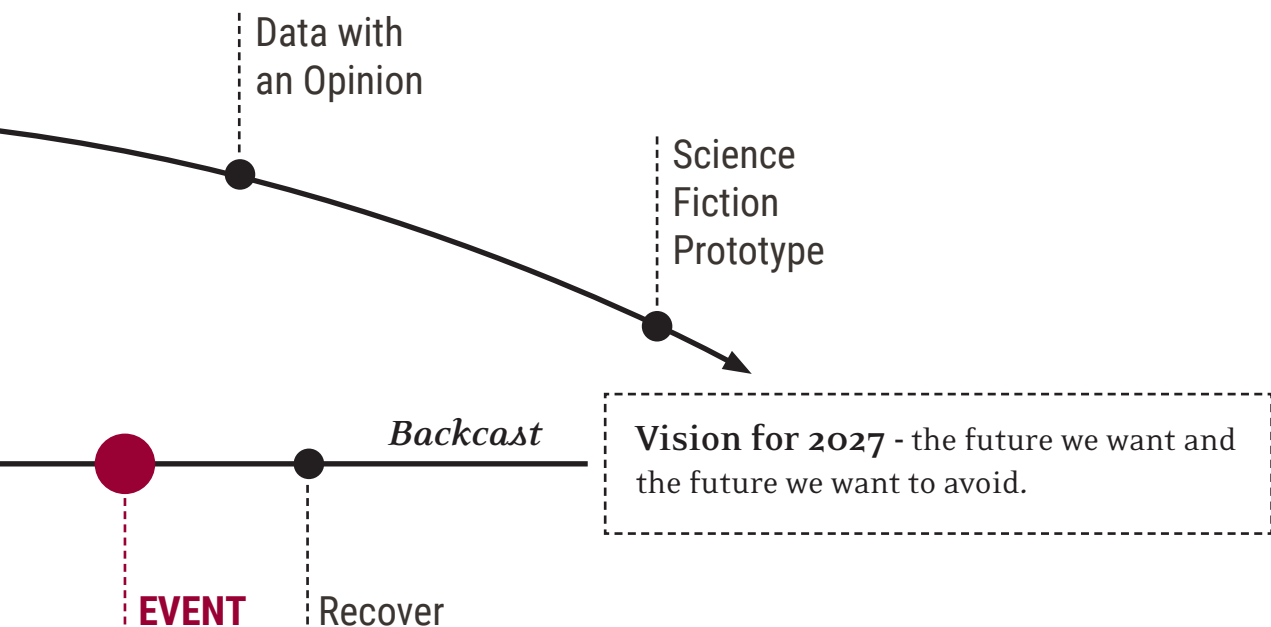
A Brief Overview

Threatcasting is a conceptual framework and process (see Figure below) that enables multidisciplinary groups to envision and plan systematically against threats ten years in the future. Groups explore how to transform the future they desire into reality while avoiding an undesired future. *The threatcasting process is described in detail in Appendix 1.*



Threatcasting uses inputs from social science, technical research, cultural history, economics, trends, expert interviews, and even a little science fiction. These various inputs allow the creation of potential futures (focused on the fiction of a person in a place doing a thing). Some of these futures are desirable while others are to be avoided. By placing the threats into a fiction story, it allows readers to imagine what needs to be done today and then three years into the future to empower or disrupt the targeted future scenario. The framework also illustrates what flags, or warning events, could appear in society that indicate the progress toward the threat future.

Threatcasting is a human-centric process, and therefore the humans that participate in a threatcasting session are important. Diversity of age, experience, and education within small groups are key but tied to a common thread - they are practitioners. Threatcasting is a theoretical exercise undertaken by practitioners with special domain knowledge of how to specifically disrupt, mitigate, and recover from theoretical threat futures. Additionally, a few participants are curated to be outliers, trained foresight professionals, and young participants for a fresh and multi-generational perspective in the groups. When using threatcasting on military problems, the mixture of participants are from academia, private industry, government, and the military.



The AI Weapons Factory

Ahmed looked across his desk at the smiling government official. As the president of the university, Ahmed had the power to do what the man was asking, what his country was asking him to do. Their government didn't have the people or expertise for the weapons they wanted, but what they did have was money.

Ahmed had been friends with Gill Dougherty since their college days in the United States. Gill was a genius even then. Ahmed had known it far before his friend became the tech millionaire he was today, the dreamer, and the visionary. He was aware that Gill needed the funding and the support to develop his most ambitious project to date, a super AI that could manage the world's energy and end climate change. Mired in politics, the US wouldn't help, but a small, wealthy, energy rich country like Ahmed's could be a perfect place for him and his team of researchers to save the planet.

Gill readily accepted when Ahmed and the university offered their support. The team of American researchers relocated and began work, trying to save the world. But they didn't know that what they were actually building the world's largest AI weapons factory, with the capability to invade every country in the region. Gill and his team were not saving the world-they were ending it.

**Their
government
didn't have
the people or
expertise for
the weapons
they wanted,
but what
they did
have was
money.**

Based on:

1. Future model 2-1214 "Revenge of the Luddites - Yub Nub" - Threatcasting Workbook 2

THREAT FUTURE 2

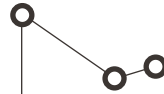
A Hole in the Heart

Harmony did not want to let her baby go. She held his lifeless body tight. Aziza was so tiny and frail - he had been sick longer than any six month old should be. Harmony was a single mother, and now her only son was gone. The AI doctors told her it was Sickle Cell disease, but none of the treatments worked. The pain had been so bad that Aziza screamed his voice away. A baby with no voice, and now he was dead.

With nothing to lose, Harmony plunges into a world of corrupt government, foreign NGOs and questionable technologies. She discovers that babies across Nigeria are dying, all with Sickle Cell disease and the AI doctors can do nothing to stop it.

Harmony uncovers a flaw in the AI doctors that everyone trusted so much - they aren't nearly as good as everyone believed them to be. The NGOs are blind to the source, and the government doesn't care about deaths of poor mothers and babies. Harmony wonders how a single, poor Nigerian woman with a hole in her heart can make a difference at all.

**The AI
doctors told
her it was
Sickle Cell
disease, but
none of the
treatments
worked.**



Threats

The Threatcasting West 2017 workshop generated dozens of possible threat futures, anticipating what life and threats could occur in 2027. By looking at major themes throughout these scenarios, as well as the raw data used to inspire them, key future threats emerged as well as possible actions to help mitigate, disrupt, or recover from possible future threats.

Threatcasting West focused on the use and misuse of Artificial Intelligence, resulting in three key threats and a collection of actions for a broad range of organizations (government, academic, industry, non-profits, and individuals). This pragmatic approach seeks to identify these threats and actions so that everyone can participate, making the necessary changes today for a better collective future. *The raw data and detailed threats can be found in the Appendices.*

THREAT 1

AI Surveillance and Coercion:

The New Dogs Of War

*Blood and destruction shall be so in use
And dreadful objects so familiar
That mothers shall but smile when they behold
Their infants quarter'd with the hands of war;
All pity choked with custom of fell deeds:
And Caesar's spirit, raging for revenge,
With Ate by his side come hot from hell,
Shall in these confines with a monarch's voice
Cry 'Havoc,' and let slip the dogs of war;
That this foul deed shall smell above the earth
With carrion men, groaning for burial.*

- Julius Caesar, William Shakespeare

The clearest and most apparent threat that emerged from the workshop raw data was a unique way in which AI could be weaponized. Surveillance and coercion are not new threats, but when conducted with the speed, power, and reach of AI the danger is newly amplified. To understand these “new dogs of war” it is helpful to understand the origin and meaning of the phrase.

At the end of Shakespeare's 1601 play Julius Caesar, Mark Antony is alone with the murdered body of Julius Caesar. In his soliloquy, he foresees a wave of violence and war racing across Italy, destroying the fabric of society. His line “Cry ‘Havoc,’ and let slip the dogs of war” refers to “havoc”: a military order that commanded and permitted soldiers to bring about chaos, pillage, and keep the spoils of victory. “Let slip the dogs of war”, can literally be seen as weapons in war fighting, but throughout history, “dogs of war” has been largely viewed as a metaphor. The dogs of war that Antony is referring to speaks to the

breakdown of law and order that holds society together, preventing the violence and destruction of war. By letting loose the dogs of war, he sees that all of the infrastructure a civil society has put in place to protect itself will crumble, unleashing the worst of humanity onto itself.

Multiple future threat models from Threatcasting West explored how using AIs to surveil and coerce would unleash a modern form of these dogs of war on individuals; forcing them to betray people's trust, dismantle corporate guidelines, local and national laws, as well as compromise national security. The growth of AI will cast a broader net of surveillance monitoring and data collecting about individuals; ultimately fusing this into a completely different form of information. The aggregate information from social media activity, media reports, medical records, security feeds (e.g., CCTV), GPS, and public records may be pulled into an expansive profile for a single person.

Although this kind of surveillance has already begun, the scale at which weaponized AI could achieve this would mean that a much larger swath of people could be surveilled, with minimal effort and activity on the part of the adversary.

Beyond simply deploying an AI weapon to surveil a targeted person, adversaries will use AI to discover victims for coercion. Imagine a criminal targeting a specific business or industry (e.g., banking) using their AI to surveil employees along with their families and any other individuals that would allow the adversary to gain information to be used for coercion. This kind of surveillance doesn't necessarily need to be targeted solely at humans either. With broad surveillance capabilities, AIs can examine all weak points in a target organization - human, technological, systemic, or a blend. A blended surveillance attack could then not only directly target employees, but also alter their behavior to cover the AI and adversary's tracks, creating an invisible entry into their target audience.

With this damaging information gathered, the adversary could use it to force a person or group of people to take action, effectively making the person an agent of the adversary. The goal of the adversary would depend on the nature of the threat actor (criminal, terrorist, state sponsored). Regardless, the weaponization of AI to surveil and coerce individuals is a powerful emerging threat. As a developing platform for psychological, physical, or systemic infiltration, AI is quickly becoming the realization of a modern dog of war, unleashing the worst of humanity and our technology onto ourselves.

THREAT 2

The AI Weapons Factory

Traditional munitions factories that manufactured weapons for warfighting have until now been easily defined, identified, and targeted. The factories developing ammunition, explosives, chemical weapons, and nuclear weapons each shared a similar profile - they were a physical place where materials were gathered to produce a desired weapon. More importantly, there were measurable indicators to track whether these factories were producing weapons or more peaceful outputs.

Experience with traditional kinetic weaponry allows organizations to understand the immediate mortal threat. AI weaponry shifts armament into a new paradigm that is more difficult to track, more subversive, integrated, and systematically impactful. This not only applies to individuals

but also to systems, governances, and operating norms on an unprecedented scale. Imagine the destruction of an entire city energy system or the turning of common household connected devices (Internet of Things) into malicious actors. As AI continues to be integrated into everyday mundane tasks as well as into core functionalities of cities and governments, the potential for rogue, turncoat, or altered AIs rises, increasing the potential for integrated AI threat actors operating behind the scenes.

The weaponization of AI presents a new challenge as we imagine the changing nature of “factories” where software instead of hardware or munitions are created. Over the last decade we have learned much about digital warfare. Cybersecurity, cyber warfighting, and cyber-crime have challenged how organizations defend and protect themselves. The coming weapons factories of AI will present a whole new host of ethical, legislative, and security issues.

In the future, how will we define and locate AI weapons factories? Especially as these “factories” are no longer solely buildings but a mix of virtual and substantially different facilities. Particularly as it shifts from a physical assembly and development model to a distributed and flexible network. Needing minimal raw materials to develop weaponry, the physical location of these factories could be anywhere and their identification from the outside, nearly impossible. Given the expanding uses for intelligent and super-intelligent AI, how will we tell the difference from a location that is manufacturing AI for the creation of weaponry versus creating AI for an innovative new gaming platform?

The regulation of these factories will also be problematic. Who is allowed to develop these weapons and where? Currently, there are no treaties, norms or laws that govern their development in the United States or on foreign soil. How can the government support technological innovation while at the same time protecting national security? How do we guard against an AI developed for constructive and lawful purposes being corrupted and weaponized for use in an entirely new way? There are precedents set in the traditional weapons manufacturing arena that can be used as a starting point. But the scale and reach of AI means that not all precedents will apply and there is significant research and work to be done.

THREAT 3

Careless Destabilization of National Security

In the last year, the emerging effects of AI on the labor force, medical industry, and culture in general have been well documented. Much of this documentation focuses on the potential for economic disruption to the workforce and the need for retraining workers with the new skills needed by the industry. Included in the core of this discussion are considerations around ethical issues for the use of AI - whether that is to decide courses of action for autonomous vehicles or to manage and treat patients. The new ways in which people will act and interact with sentient technologies are still nascent, and have emerged as a top priority in the discussion on the impact of AI on larger society.

Each new innovation and technology brings promise, but also introduces new threats. These threats are brought about by vulnerabilities introduced through the introduction of AI into systems changing the operating norms and creating new methods of exploitation. Chief among these threats regarding

the development of AI and its effects on society exists the potential destabilization of the United States and world economy, the loss of trust and acceptance of AI, and a cultural backlash against the use of AI in general. Each of these threats has the potential to be a massive destabilizer. This destabilization will certainly have adverse effects on specific nations, economies, and industries, but it will also pose a direct threat to National Security.

“Careless destabilization” is the commonly accepted idea that AI most likely will have destabilizing effects across multiple areas of society including economic, social and cultural areas. As of yet, these areas have only been partially defined. If we continue to simply track possible AI threats (e.g., embedded ethics, inherent bias, and decision frameworks), but do nothing to correct the course of development and deployment, we are being careless. For both benefits and threats, the integration of AI will impact the vast majority of American citizens, changing how they interact, work, and live. And this impact is not restricted within U.S. borders but will have impacts on societies and citizens all over the world.

Although clearly more research is needed, it is imperative to take immediate pragmatic steps to lessen the destabilizing impacts of nefarious AI actors. If we are better able to understand and articulate possible threats and their impacts to the American population, economy, and livelihood, then we can begin to guard against them while crafting a counter-narrative. How can we envision a future where AI in the workforce benefits individuals, organizations and the nation? How will medical AI allow us to live longer and make the healthcare industry more manageable?

Currently, there is a race toward creating the first “true” AI. Because the stakes are so high for the disruptive impacts of AI on all aspects of our society and culture, it is imperative that America is the leader in fully developing and implementing AI. AI might not only integrate with systemic operations across civic and industrial organizations, but could also integrate into the very way citizens behave. Imagine Siri, Alexa, or Google Home shifting toward subtle behavioral nudging of millions of users based on an adversary hack. Now think of a truly sentient AI and its behavioral modification capabilities. This can bloom into a national security problem. An entity (nation-state, organization, or company) that masters “true” AI functionality even a week before us will have a significant technological advantage over the U.S. While it’s comforting to imagine that the U.S. or a friendly nation-state achieves victory in creating “true” AI, we have to assume that it won’t be. This game-changing technology has a steep learning curve - and the consequences of not keeping up are potentially devastating.



THREAT FUTURE 3

The Edge of Nothing

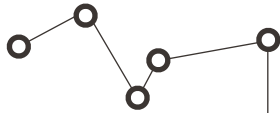
Harriet Downs had it all: a great job, a loving husband, and two beautiful children. She was an up-and-coming programmer at Goldman Sachs for the company's essential AI trading bots, on the fast track to management. On this day, Steve and the kids were getting settled into their beautiful new house in Sevenoaks while she took the train into London. Harriet Downs had it all until that day on the train when the man with the lion tattoo on his neck stopped her and showed her the video.

She recognized the people on the screen. One of them was her. She remembered the terrible mistake she had made that night. Too much to drink. Too much stress at work. It was never going to happen again. But somehow the man had gotten a video, knew everything about her life, her habits, her family, her work. And he wanted something...

It was a simple piece of code that needed to be inserted into the bots at work. No one would know or understand why the AI was selling millions of shares at once. Yes, the markets would collapse but just for a moment, just for a second, just long enough for the man's "friends" to make billions by shorting the stocks.

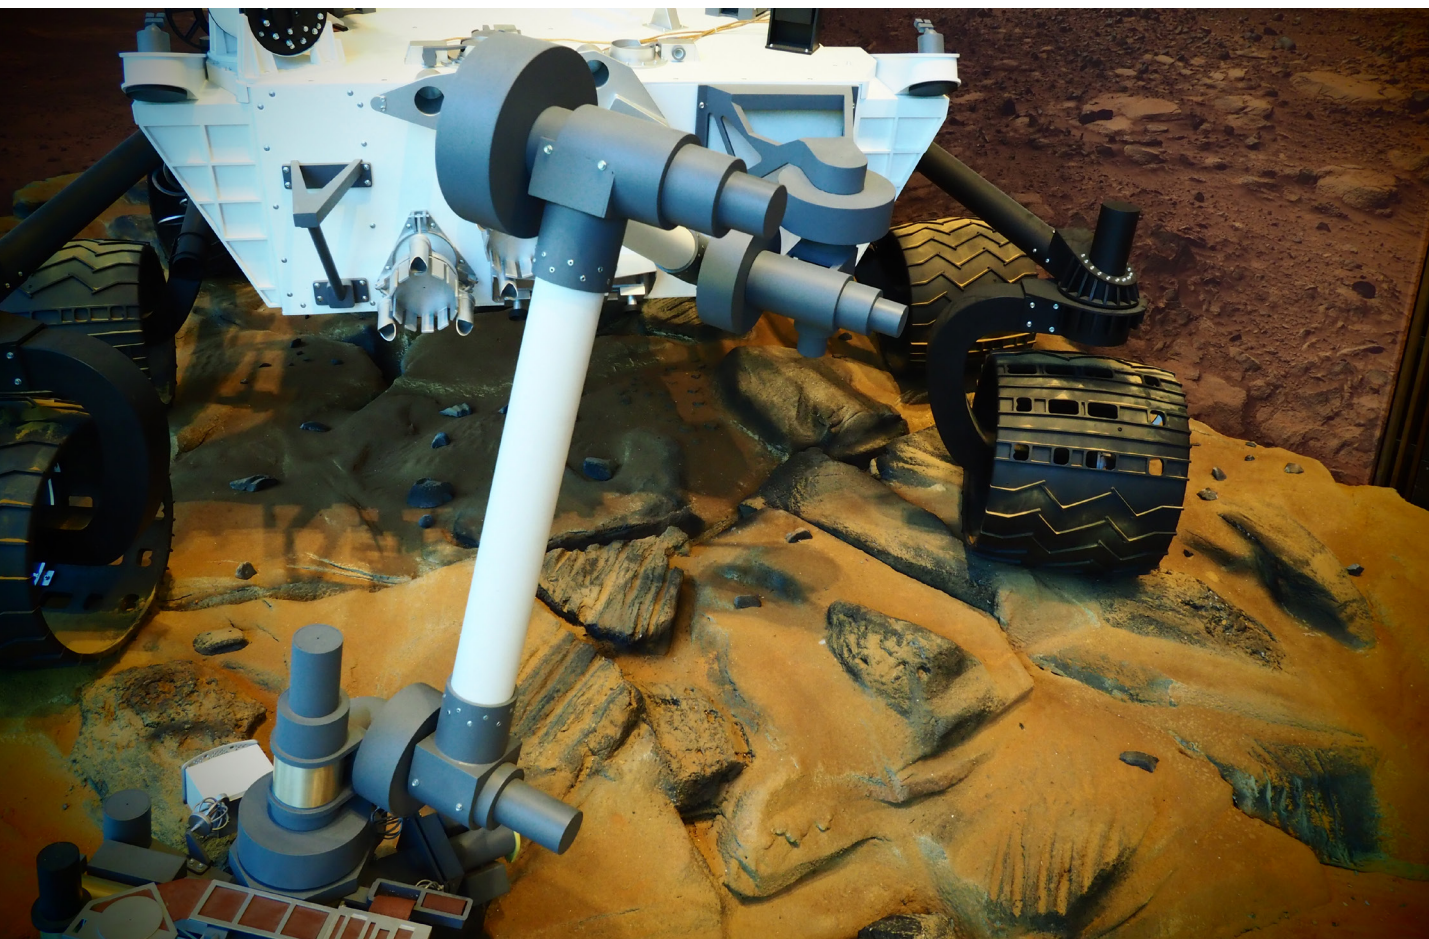
Standing on the train, Harriet had a decision to make. Did she want to have her life back or be the woman who had nothing? No family. No job. Nothing. But what Harriet didn't know was that there were other AIs watching her as well, they knew she was at risk and they just might be able to help her step back from the edge of nothing.

It was a simple piece of code that needed to be inserted into the bots at work.



Actions

The Threatcasting West 2017 workshop uncovered not only threats but also actions that could be taken to help mitigate, disrupt, and/or recover from the threats. Three high-level actions are centered on further research, technological and process tools, and regulation and oversight. These actions constitute a “whole of society” approach to problem solving and have been applied to specific domain areas with detailed steps that can be taken. *Raw Data and details can be found in Appendix 4.*



ACADEMIA

Academic research is critical to the economic and social development of society. Sometimes the results of research are easily apparent, but at other times, the benefits are not initially obvious. According to Einstein, “If we knew what it was we were doing, it would not be called research, would it?”

Many of the suggested actions to take against our futures require research, exploration and innovation that our students and faculty at higher education institutions in this country could explore.

- Develop academic programs, courses, concepts and content that include ethical behavior when thinking about the development of AI and algorithms.
- Incorporate into research the implications of AI becoming highly developed and its impact on the future workforce.
- Conduct research focused on creating an AI that can evaluate decisions, monitor ethical practices in other AI systems and remain ethically compliant in its actions and decisions.

23

INDUSTRY

Industry plays a valuable role in the solution space to mitigate future societal threats. Industries should consider creating greater open source environments where others within the industry and outside of it (e.g., governments, non-profits) can learn.

- Develop algorithms that have a system of checks and balances built within themselves. An example: Each algorithm could have a cluster of algorithms associated with it - an ethics algorithm, a social algorithm (looks at how the main algorithm is affecting the humanity component), and a detective algorithm to name a few.
- Consider implementing “kill” switches in AI which use a mechanism (digital or physical) that temporarily disables or locks the AI without destroying it completely.
- Explore greater implications and algorithm scenarios, develop negative scenarios, explore how the algorithms or AI actions could present themselves, then develop a new scenario with the same algorithm, illustrating how it can be used for good.

THREAT FUTURE 4

Little Sister Lost

Ba Wei was the smartest girl in her class. She did what her parents expected, and in her neighborhood in China, she was a model citizen. Even with tensions rising between the United States and China, Ba, her younger sister Ju, and their parents secured a visa for both girls to study engineering in Portland, OR.

When the family arrives in America and the girls enter Portland State University, things begin to go terribly wrong. Ju pulls away from the family and tragically takes her life, throwing herself off the roof of a parking garage. Distraught, Ba begins friendships with new American friends. It's then she notices that something is not right...

Ba's social media accounts and online profiles begin acting strangely. She notices the AI is behaving erratically and believes something or someone is actively trying to keep her from her new friends, deleting texts, disabling phone contacts. Always the good student, Ba learns that she's being manipulated by a Chinese program to control its citizens on foreign soil and worse than that - her parents are complicit.

Shocked, Ba begins to question... was she manipulated too? What else have the AIs been up to? Has she been manipulated her entire life? Was this what pushed Ju to take her own life?

Ba learns that she's being manipulated by a Chinese program to control its citizens on foreign soil and worst than that – her parent are complicit.

Based on:

4. Future model 1-12 "Let's Dance" - Threatcasting Workbook
5. <http://www.cccb.org/en/multimedia/videos/the-state-of-surveillance-big-brother-little-sister-and-uncle-sam/211473>

INDUSTRY CONT.

- Ensure that no one person has too much authority to datasets, or data warehouses. Implement whistleblower programs to prevent people from being blackmailed.
- Explore methods to influence AI including nuanced information operations in the wake of nefarious AI actors. Industry might need to change their business best practices from static reaction to a dynamic response.
- Explore greater implications and algorithm scenarios, developing negative scenarios, exploring how the algorithms or AI actions could present themselves then develop a new scenario with the same algorithm, illustrating how it can be used for good.

NON-PROFITS

Non-profit organizations are created to champion and enrich the well-being of our communities. They provide the opportunity for people to join together and contribute their time, resources and experiences to serve a greater good. These organizations have a vital role to play across the action space.

- Advocate for developing national legislation that outlines data protection measures to preserve privacy and integrity of data associated with US citizens.
- Encourage industry organizations to develop standards and guidelines that support data integrity and security within the development of new digital technologies rather than as an afterthought.
- Inform the customer about the security of digital technologies that they bring into their home and family. Help build the framework that the Cyber Independent Testing Lab (CITL) is developing.
- Become a champion for the general public's measured and pragmatic understanding of AI.
- Develop trainings and materials to better inform and equip the industrial workforce for working securely with AI.

INDIVIDUALS

We each have a responsibility to be informed and protect our personal data. As we become more and more comfortable with intelligent digital systems, we will naturally start to innately rely on and trust them. Our biases will change and evolve over the next few years, and a healthy dose of awareness will go a long way.

- Question how your personal data is being used and the implications, both positive and negative, of sharing data.
- Trust your gut. Don't trust blindly. If something seems wrong, it very well may be.
- Demand that brands and organizations practice transparency and inform you of how they are using your data.
- Champion awareness with populations and communities without access to training or education about AI safety.

GOVERNMENT

While there are actions that individuals and organizations can take to help protect our future, there are also actions that clearly fall into the realm of government's responsibility. Both AI and human checks and balances are needed.

- Explore and debate the development of international disclosure laws. Example: if a country, organization or other threat actor contacts an organization or individual about using AI for harm, there is a requirement to disclose it. This will create both interior and exterior transparency.
- Explore and debate if specific computer hardware that enables AI should be registered with government or another organization. Develop requirements for how hardware is positioned and how it can be used. Government organizations, like the FDA, could certify suppliers. Prototype frequency of messaging and amount of data being given to any consumer at one time.
- Explore the creation of an international organization that can oversee the development of AI to ensure that it is not weaponized.
- Keeping humans in the loop of processes will be imperative. Following are some examples:
 - Human intelligence (HUMINT) will continue to be needed. Humans may be the primary methodology for differentiating between real people and digital avatars.
 - At ports and other entry points that rely upon AI for security, have spot inspections conducted by humans. Distribute the AI by requiring separate AIs for inspecting the goods, receiving those goods, and for security.
 - Design backup systems for vetting employee data that are human-controlled and regularly checked. ○

THREAT FUTURE 5

Trust Fall

When Dr. Lei Lin was a little girl, all she wanted to do was save the stray cats of Taipei. She spent her life worrying about the ones no one else cared about. It didn't surprise anyone in her family when, after becoming a doctor, she moved to rural Taiwan to serve the poorest villages. With little funding and no support staff, Dr. Lei learned to rely on her medical AI to keep her current and solve the unique and tricky problems of her patients. It was not easy work, but she knew it was important.

On a rainy Wednesday afternoon, Dr. Lei was checking in on a young mother who had been experiencing symptoms of GI disease. Typically she would prescribe hyoscyamine or dicyclomine depending on which was easier to get. Today, however, her AI recommended something different. For the first time in the two years working in the Taipei community, she disregarded the AI and relied on her own training.

The AI had been wrong. It was quite clearly wrong. Had it been wrong before?

By the time Dr. Lei returned to her meager practice, the minor incident had begun to weigh heavy on her mind. The AI had been wrong. It was quite clearly wrong. Had it been wrong before? Had she just not noticed it? Why was it wrong for a medication that was so simple? What was behind it? More importantly, who was behind it?

Based on:

6. Future model 1-93 "Hacked Doctor" - Threatcasting Workbook 1



Post-Analysis Themes

The post-analysis of the raw data and threat futures conducted by the Threatcasting Lab synthesized and clustered a group of compelling themes that touched on and ran throughout many of the threats – long processes and preparation, trust, job loss, human as the hero, and a need for new norms.

Complete raw data and threat details can be found in Appendix 3

THE LONG GAME

In the development of these futures, many of the threat actors took a low and slow, long game approach. Meaning, it took a significant amount of time to achieve the threat objectives as preparation of the battlespace and actions towards the targets were lengthy in process. There were no short term weapons or threats like “wifi guns” or “data grenades”. Instead, threat actors had to look towards the long game to support their actions. Which leads to the question, in 2027, are quick attacks in the cyber domain no longer possible or probable? Given a world with compute power everywhere and sensors surrounding us, do malicious behaviors have to be so hidden that only the low, slow game can be successful?

TRUST

Another recurrent theme in the futures centered on trust. Given today’s behavior and extrapolating to tomorrow, people innately trust digital technology. If Siri tells us an answer, it is assumed by many that it must be true. If Google provides a search result, it is also assumed that it is a fact. If something is re-tweeted hundreds of times, it must be real.

It seems that unlike other technological advances in human history, digital technology advances begin in a state of trust by the general public. For example, most people didn’t feel safe flying until after World War II. This immediate and complete trust in digital technology is disconcerting and troublesome because if a threat actor disrupts this innate trust, the trust in that technology could be gone forever. This could lead to people losing trust across a broad spectrum of technologies and once this trust is lost it will be difficult to recover.

How do we instill a societal behavior to trust AI because it could bring great good while at the same time not giving it blind trust to enable resiliency? We need to question digital technology when the results do not make sense. Fake news has been around for hundreds of years. However, digital news and recent political issues have spotlighted the algorithms that populate news feeds and the neglect these algorithms have for truthfulness or objectivity. For many, any news that does not support one’s biases is deemed fake. The threat futures generated by the workshop struggled to develop a balance of trust and questioning with digital technology.

As a society, how do we take basic precautionary measures to protect ourselves from malicious behavior but not to a point where decisions are based solely on fear? A good example in the physical world of this measured approach would be locking one’s doors at night, but not being so frightened that one barricades one’s family inside the home with an arsenal of weapons for protection. The goal should be to achieve the maximum economic benefit AI offers while being aware of the risk that new technology brings.

JOB LOSS

The looming specter of technological unemployment and job loss could be found throughout the threat futures of 2027. During technological revolutions in human history, the then-current workforce structure evolved. Old jobs were lost, new jobs were created, and a training/education plan was crafted to support this transition. Many of the future visions illustrated threat actors taking advantage of societal destabilization due to workforce changes.

There is an urgent need to have a truthful conversation about AI, the job loss potential, and the destabilization effect it could have on society. It is clear that the adoption of more automation within society will result in the loss of work roles. But there is also a current alarming conversation perpetuated by recent media attention and for-profit re-training businesses that could have a greater destabilizing effect on society than the actual job loss. The escalation of this issue without a measured debate is bad for both national and international security. It could lead to a wide-scale public distrust of AI, preventing US businesses and research institutions from harnessing the benefits and adoption of innovation. This could leave the nation behind other nations and economies who have accepted and continued their development of these technologies. To disrupt this narrative, a balanced counter-narrative could focus on the potential evolution of education and training that will be needed to deal with this next chapter in human history.

HUMAN-CENTRIC

In most of the threat futures, humans were the heroes of the fictional accounts. Threat adversary advances made through our own and trusted digital technology were almost always detected by a human. The futures shared the idea that a human actor (outside the process) detected symptoms of the low and slow malicious actor's plan and investigated them. Once the attack was successful, they questioned how the situation came into being and connected the dots. If a human could do this post-attack, how could an evolved-AI not detect this

pre-attack before damage is done or worse lives are lost? Is there an ability to create a technological problem detector? How do you protect it from being corrupted? How do we meaningfully create the digital equivalent of a private investigator, investigative journalist or oversight committee? Can we create something that humans would trust?

A strong undercurrent in the threat futures was a need for a new normative behavior focused on digital technology security. In a future where AI is surveilling and coercing, what are the backstops that we can place in our environment to alert humans to problems and keep them safe? How do we create these new norms?

These backstops may not deter everyone but can be helpful to most as we construct our relationship with evolving digital technology over the next decade. How do we create a certain amount of protection that would deter and protect most people, much like the posted danger signs, extra fences, and patrolling police on the beach of Saint Martin, that attempt to warn off individuals from standing in the path of jumbo jet engines and from being violently blown into the water. While nothing will stop die-hard thrill-seekers, it can protect unsuspecting individuals from injury.

As AI is developed and increasingly deployed, what can government, industry, and education do to help? Much like current laws that prevent texting while driving, how do we create new norms for behavior in the cyber world to protect people? How do we translate these normative rules from the physical domain to the digital domain?



THREAT FUTURE 6

Old Dogs. New Tricks.

With the
clock
ticking,
Major Moore
digs in...

The Department of Defense's multi-force engagement in the Middle East was set to begin in 13 days. Across the nation, and the globe, United States forces rush to action. At Hill Air Force Base Major Don Moore put in long hours monitoring the essential equipment and vital support needed for success. Intricate supply chains stretched around the planet and on this night they were all focused on the same place, but something was wrong...

With the clock ticking, Major Moore digs in to find a shocking player behind these little inconsistencies. Hot on a data trail, he watches what he thought was a blip expand before his eyes into a massive surveillance web throughout the supply chain. An AI is watching every move, recording decisions, and reporting back to its master: Russia.

Major Moore knows that by sounding the alarm, he will disrupt the pending mission, disable the AIs, and throw the system into chaos. Isn't that what the Russians want? A small advantage to gain supremacy in the region? Before he can act Major Moore discovers something that makes his blood go cold - they have been watching for 10 years.

Based on:

7. Future model 2-1 "The Perfect Red Team" - Threatcasting Workbook 2



Current Work and Implications

In the post-analysis, looking outside the Threatcasting West 2017 workshop data to current work and implications uncovered signals and evidence in support of the threat futures, with differing perspectives, industries, and applications. Using this data and research, it's possible to map how these threats might emerge, and in what ways they are already being realized. *The examples provided are not intended to be a complete exhaustive survey. They are provided for further perspectives.*

AI SURVEILLANCE AND COERCION: THE NEW DOGS OF WAR

Over the past few years, the weaponization of AI and the future effect it could have on society has become a topic of debate. Organizations like the IEEE's Global Initiative for Ethical Considerations in Artificial Intelligence & Autonomous Systems and the Centre for the Study of Existential Risk at the University of Cambridge have gotten involved, and the conversation has made it all the way to social media via the "Twitter war" between billionaires Elon Musk and Mark Zuckerberg. AI weaponization has become a matter for serious discussion as well as media hype.

THE SUPER-AI DEBATE

What happens when we develop machines that are smarter than humans and can learn faster than we can? This question crops up in some of the foundational minds of the AI debate. Much of this work explores a dystopian future where our machines are no longer in our control. Such prominent people as physicist Stephen Hawking and entrepreneur Elon Musk have cautioned that the unchecked advancement in AI could be even more impactful than the splitting of the atom and the development of nuclear weapons to the human race.

"What we really need to do is make sure that life continues into the future. [...] It's best to try to prevent a negative circumstance from occurring than to wait for it to occur and then be reactive." - Elon Musk

CURRENT WORK

In spring of 2015, Future of Life Institute launched the AI Safety Research program, funded primarily by a generous donation from Elon Musk. By fall of that year, 37 researchers and institutions had received over \$2 million in funding to begin various projects that will help ensure artificial intelligence will remain safe and beneficial. Now with research and publications in full swing, we want to highlight the great work the AI safety researchers have accomplished, which includes 45 scientific publications and a host of conference events. - futureoflife.org

While research is ongoing, the question remains: what happens when our intelligence is surpassed by something we have created? Super-AI is a very real possibility and offers seeds of warning as to how we might best create guardrails to protect ourselves, particularly if such advanced intelligence can be manipulated - or do the manipulating - against a particular group.

CURRENT WORK

Humans, Not Robots, Are the Real Reason Artificial Intelligence Is Scary Intelligent weapons are too easily converted by software engineers into indiscriminate killing machines. - By Zach Musgrave and Bryan W. Roberts, Aug 14, 2015

Facebook Shuts Down AI System After Bots Create Language Humans Can't Understand - Indo-Asian News Service, 31 July 2017

The 'creepy Facebook AI' story that captivated the media - By Chris Baraniuk

Elon Musk and Mark Zuckerberg are waging a war of words over the future of AI - By Jacqui Frank, Kara Chin and Joe Ciolli

THE LETHAL AUTONOMOUS SYSTEM DEBATE

The second area of debate focuses on the use of AIs in lethal autonomous systems and the ethics of these machines. What will it be like to have robots that can kill without a human in the loop, making the crucial kill, or no kill decision? If we can no longer control the decision process for AIs that carry lethal payloads, how can we ensure the intention and action of the very machines we have built? Ethics boards in the European Union, trade associations, and universities have begun to debate this topic.

THE SILENCE AROUND MID-LEVEL AI WEAPONS:

What is not being researched and discussed are the class of AI weapons that sit just below these high profile and media-grabbing weapons. This simpler and much more possible class of weapons needs additional research and debate.

A review of current peer-reviewed and academic publications shows that little writing or scholarly study has been conducted with a focus on these possibly nearer-term weapons. It is likely that while academics are not focusing on this subject, criminal organizations or state sponsored adversaries might use the ability of AIs to surveil individuals and groups inside companies to attain subversive goals. A history of utilizing emerging social media technologies in novel ways (such as Twitter use by ISIS) points toward the likely eventuality that the emerging technology of AI will be used to circumvent traditional channels of communication or espionage.

TECHNICAL INDICATORS: IS IT POSSIBLE?

A review of technical publications and peer reviewed journals show that this class of AI weapon to surveil and coerce is nearly possible. Extensive research is being done to use AI to track the spread and response to medical and social epidemics (Innovations in Population Health Surveillance: Using Electronic Health Records for Chronic Disease Surveillance and Collective activities in a technology-mediated medical team. An analysis of epidemiological alert management), as well as using AI to monitor social network activity for terrorist threats and malicious actors (The Rise of Social Bots).

From this review, it is apparent that weaponized AI to perform surveillance across a wide array of inputs is possible. It appears that in most instances a human actor would use this information to coerce an individual to take action against their interest. But with the rise of attacks like ransomware, one could see how an AI all on its own could gather information on an individual, contact that person, and get them to take action leveraging personal information as a controlling device, and asking for subversive action rather than monetary payment.

Prominent media distributors are now looking at what it might take to manage the growth of an AI arms race, as many observe it has already begun. The global community is now coming to terms with the idea that AI weaponry is on the near horizon, so how might we best prepare? The article “It’s already too late to stop the AI arms race—We must manage it instead” takes a deeper look.

CURRENT WORK

Can we prevent an artificial-intelligence (AI) arms race? While an ongoing campaign argues that an agreement to ban autonomous weapons can forestall AI from becoming the next domain of military competition, due to the historical connection between artificial-intelligence research and defense applications, an AI arms race is already well under way. Furthermore, the AI weapons challenge extends far beyond autonomous systems, as some of the riskiest military applications of artificial intelligence do not select and engage their own targets. This article draws on the history of AI weaponization and arms control for other technologies to argue that artificial-intelligence and robotics researchers should cultivate a security culture to help manage the AI arms race. By monitoring ongoing developments in AI weapons technology and building the basis for informal “Track II” diplomacy, AI practitioners can begin building the foundation for future arms-control agreements. – “It’s already too late to stop the AI arms race—We must manage it instead”. Edward Moore Geist. Bulletin of the Atomic Scientists Volume 72 Issue 5. 9/2/2016.

THE AI WEAPONS FACTORY

When considering the future of AI weapons factories, we must first understand that this kind of factory is fundamentally different from any other situation we have dealt with in the past. Historically, it has been possible to identify a weapons factory by particular characteristics such as raw materials, specific equipment, and personnel. For example, nuclear weapons factories require plutonium, centrifuges, and physicists. When intelligence indicates that an adversary is collecting and amassing these elements, that intelligence gives us an early indication that the adversary is building a nuclear weapons factory. These indicators are nuanced enough that they enable the distinction between a factory crafting nuclear warheads and a country constructing a sustainable power plant.

A LACK OF DEFINED INDICATORS

These kinds of specific indicators are yet unenumerated and untracked as it relates to possible AI weapons factories. To complicate matters, the possible elements of an AI weapons factory can also comprise the raw materials necessary for a computer game development studio or a financial technology trading company.

The necessary elements for weaponization and simple consumer electronics or technology development means that not only could false positives be identified (e.g., a small group of entrepreneurs or innovators as adversaries) but it also adversaries could use seemingly normal industrial or research activities to hide AI weapons factories. Before the process of regulation or treaties for AI weapons factories can begin there will need to be an accurate and defined means to be able to identify and validate specific indicators.

It took decades and international treaties to determine the best way to monitor the potential for the weaponization of nuclear technology. This understanding was aided by the fact that the resourcing required for nuclear weaponry creation required sourcing materials amongst nation states. Even then, global understanding and creation of treatises took time and

financing. The digital realm does not have the same barriers to entry for the potential weaponization of AI. As global connectivity increases, access to digital resources doesn't typically depend on trade between nation states.

TECHNICAL INDICATORS: IS IT POSSIBLE?

The Google Loon project is focused on allowing internet access for everyone - independent of geographic location. Given the rise of desire for interconnectivity, and easy access to tutorials (ranging from simple to complex) it is not a stretch to imagine organizations or individuals gathering the necessary information to build semi-intelligent machines - particularly with the trend toward open source data sets (Google and Facebook have already opened their machine learning and AI datasets to the public). These trends, along with the rise of cyber warfare, indicate that the weaponization of AI is on the horizon; and to use AI weaponry at scale, manufacturing outposts will be a key component.

LACK OF RESEARCH

There is a significant lack of academic research and thinking about what potential indicators might look like within the digital and physical domain to ensure a minimization of false positives while continuing to support progressive innovation. How would you track the brainpower of key scientists and researchers as they shifted from game design to country overthrow? What actions in the cyber domain would correspond to an action in the physical environment? Is it even possible to monitor for the weaponization of this technology in the future or do we need to develop a new strategy to secure our society/citizens/way of life?

The Automatic Spy

**They are
afraid that the
government's
pervasive
information
systems – now
linked through
AI – gave the
government
unprecedented
access to
Americans'
personal lives.**

This was desperation. Rahul stood before the lavatory sink ignoring his reflection in the mirror. What choice did he have? It wasn't even really a choice. A choice between anything and self-destruction was no choice. Everything he'd worked to achieve, gone in an instant. They were right: the Air Force wouldn't just revoke his clearance, more likely they'd find him guilty of willfully misrepresenting his medical history. Even if he weren't court-martialed, he'd be put out of the service with nothing. No pension. No disability. No resume--who would hire a disgraced officer?

In the media, they seemed almost noble. Anti-tech activists, the journalists called them. They are worried about the intrusion of artificial intelligence into daily life. They are afraid that the government's pervasive information systems - now linked through AI - gave the government unprecedented access to Americans' personal lives. And they weren't asking him to take anything or even tell them anything. Delete their information from the systems. He had the access and the ability to manipulate the logs; he knew he wouldn't be caught. That wasn't what worried him. No one else would know, but he would know. And then he again remembered the medical files - clear proof he should hold neither the clearance nor the position he now enjoyed. What kind of choice was it?

Rahul looked himself in the mirror then quickly stepped through the door and strode purposefully back to his workstation. His mind was made up.

CURRENT WORK

“Fragile and conflict-affected states incubate and spawn infectious disease, illicit weapons and drug smugglers, and destabilizing refugee flows. Too often, failures in governance and endemic corruption hold back the potential of rising regions. The danger of disruptive and even destructive cyber-attack is growing, and the risk of another global economic slowdown remains”- 2015 National Security Strategy

The United States national strategies include responses to concerns about when and how other nation states may devolve into a failed or failing state. However, we should not blind ourselves to only imagining this condition to happen to others in the future, and we should not only consider physical destabilization but digital destabilization as well.

Since World War II, the United States has acted as the primary force to maintain international security and stability, leading first the West in the Cold War confrontation with the Soviet Union and, more recently, international efforts to confront violent extremism. Driving these efforts has been a set of enduring national interests and a vision of opportunity and prosperity for the future. U.S. interests include protecting the nation and our allies from attack or coercion, promoting international security to reduce conflict and foster economic growth, and securing the global commons and with them access to world markets and resources.

CARELESS DESTABILIZATION OF NATIONAL SECURITY

Destabilization is well-known as a tactic for weakening traditional organizations. From economic destabilization to terror attacks, adversaries try to instigate or capitalize on the destabilization (of troops, communication, sustenance, etc.) of an enemy to gain strategic advantage. Each destabilization can ultimately threaten the national security of the impacted country.

DESTABILIZATION OF TRUST

Several of the threats modeled in this report ultimately brought about a destabilization of trust for the United States including a decrease of trust in the economy, trust of the accuracy of technology, or trust that a citizen has in their government and societal structures. Destabilizations like these may be brought about by job loss, economic downturn, and failure of medical infrastructures to adequately give care.

Current conversations regarding the potentially destabilizing effects of rampant AI represent a bleak view of the future. A massive terror attack is not the only way to undermine the normal functioning of American society. A threat actor could target the destabilization of trust, nudging, swaying and ultimately convincing citizens that we are socially or biologically doomed and that the government will be of no help.

CURRENT WORK

The security of the United States is tightly bound up with the security of the broader international system. As a result, our strategy seeks to build the capacity of fragile or vulnerable partners to withstand internal threats and external aggression while improving the capacity of the international system itself to withstand the challenge posed by rogue states and would-be hegemon.

- 2008 National Defense Strategy

CURRENT EXAMPLES

Destabilization events have already begun to occur on smaller scales. A leak of personal account information from the popular adultery site Ashley Madison created social waves in 2015. Yahoo closing its doors after plummeting stock performance left many questioning the viability of digital mega corps and the future of the digital economy. Both of these incidents happened following cyber-attacks, and while these examples are not yet on the scale we are talking about, they are hints of what is to come.

Throughout 2017, ongoing cyber-attacks targeting businesses across Europe and Ukraine have had massive economic destabilizing effects. While these attacks were not AI-lead, they are an early indicator and potential precursor of the breadth and damage possible from AI integration into systemic control of business operations. These destabilizations, while widely felt, were but early signals of a much larger and more effective destabilization that could be brought about with weaponized AI.

THE TRUE AI ARMS RACE

What if the first nation to reach “true” AI does not share the same moral or ethical framework as us? What if their definition of the “common good” is not similar to ours? If a country opposite in moral and social norms is first-to-market, they will control the supply of AI which could result in AI applications or decision-making that is unacceptable to our society. Do we then prohibit the use of these technologies within U.S. borders until we can develop a branch of the technology that works within our societal framework? Such a shutdown could potentially result in a loss of market share and US competitiveness within the world economy.

Considering all these factors, does the United States need to start using the technology and attempt to create a patch to provide a level of security, safety, and moral responsibility onto the AI so that our society feels comfortable using it? Given product adoption history, this tactic never works well - least of all elegantly - and the potential backlash from society when the technology does not perform as desired could be devastating. The only solution left: we need to be first to the AI market.

The United States has a long history of assisting other nation states when they are troubled. Historically, this has been an effort to provide stability and sustain peace - a truism that the U.S. Army has lived throughout its 242-year history. Over the last three decades, we have intervened in several small failing states (i.e. Haiti, Somalia, Balkans, Libya, Syria to name a few) resulting in mixed results with global implications. The global impacts of a world power failing are much larger, and as of yet incalculable.

CROSS BORDER DESTABILIZATION

History shows us, however, that instability starts small and locally. When normal societal functioning begins to be undermined, the ability of government to provide a safe and secure environment for its citizens is likewise hindered. Viewing small state destabilizations as precursors to larger global destabilizations provides a warning of emerging trends. Planning is needed for possible destabilizers stemming from social, economic, and digital actions which could cause ripple effects across countries of power.

THE NEED FOR A NEW NARRATIVE

Is it possible to anticipate and counteract destabilizers from such a wide possible range of sources? To counteract these destabilizing effects a pragmatic, balanced, and well-informed counter-narrative is needed before AI's use becomes ubiquitous. Currently, this is missing from media, academic, and conversational coverage.

The expertise and research is available to craft understandable narratives of how AI is a benefit, and why organizations and the general public should design in security from the inception rather than an afterthought. Investment is needed in the development of "true" AI to ensure culturally appropriate values and compliant ethics are embedded into the first-to-market offering. Support networks can be created to help stabilize potentially unstable global interactions. Initiatives such as these need to be prioritized across sectors and industries to preemptively set expectations for the integration of AI into our world rather than trying to disseminate recovery transmissions after a destabilizing event has occurred. ○

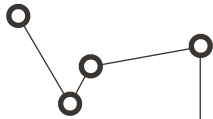
James and the Giant Truck

“Made redundant.” The phrase rattled James’ thoughts as he rode along, a British sounding euphemism - made sense, as he’d read it in a BBC article. In that last six months, two-thirds of the city’s sanitation workers had lost their jobs, and they all used the American term: fired. The City of Houston had become the third city globally to adopt fully automated waste management. Self-driving trash trucks are what the drivers called them.

Fifty-five years old and in six months he too would probably be let go. Hundreds of hours of testing in the streets of Houston had demonstrated that his presence wasn’t necessary. The machines didn’t make mistakes. They stopped when pedestrians stepped in their path. They went from highway to avenue to side streets without interfering with traffic. They emptied dumpsters that weren’t even where they were supposed to be. The only reason he still had a job - riding along, babysitting the machine, ready to press a stop button if something did go wrong - had been a political concession to the union to save jobs for a few workers like him. Waste management indeed, he thought. The drivers were the waste to be managed.

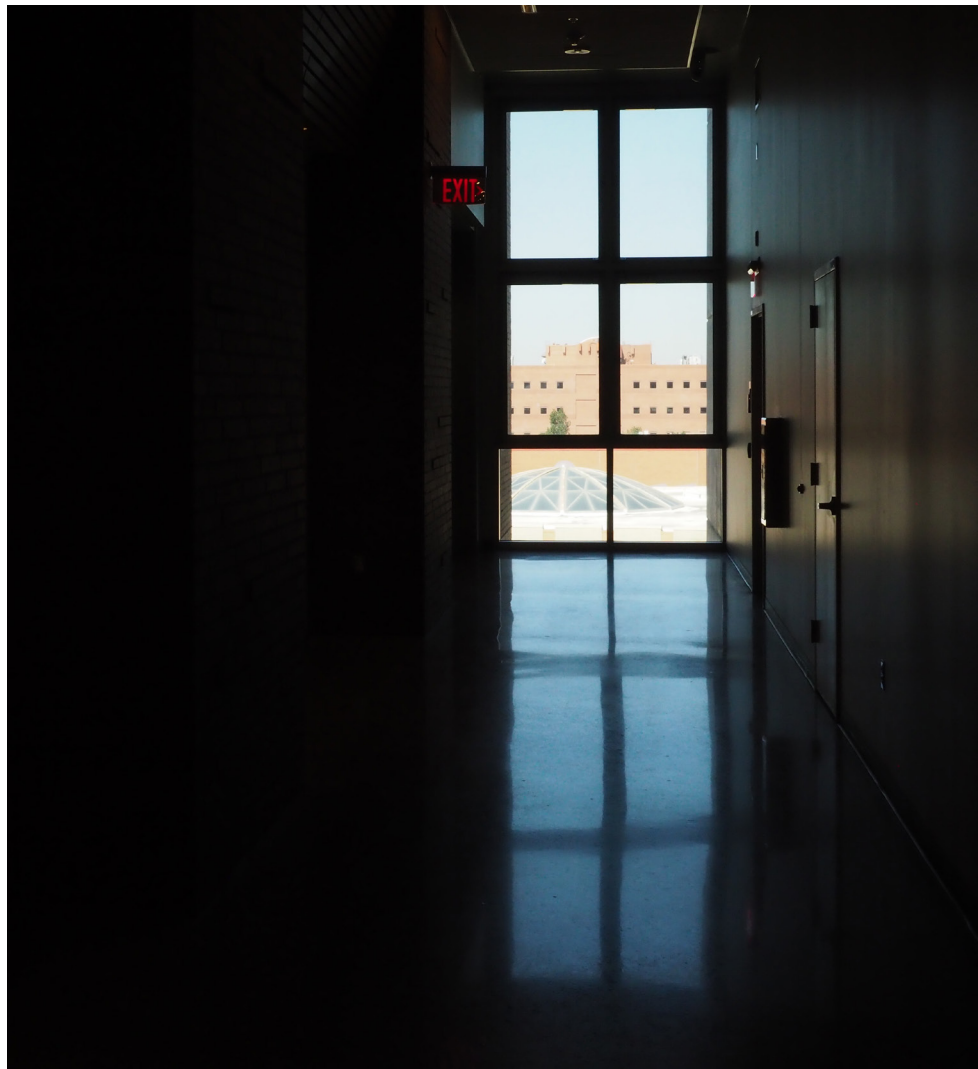
Then, last night he’d read that BBC article about the city of London. Then, last night, another driver shared the BBC article in their private Facebook group. The City of London had been forced to rehire drivers after a number of safety incidents. No one could explain the problems. More than a few thought the drivers’ union was behind it, but no one could prove it. Now as he rode, mentally turning the phrase “made redundant,” James thought about what it would take to make the system fail.

**Hundreds of
hours of testing
in the streets
of Houston had
demonstrated
that his
presence wasn’t
necessary.**



Conclusion

In 1911, Ernest Rutherford (the “father of nuclear physics”) could not have imagined that his ground-breaking, scientific research into understanding the atom would lead to Hiroshima and Nagasaki over three decades later. Just as Arthur Samuel (the “father of artificial intelligence”) in 1959 could not have imagined that teaching an IBM mainframe to win at checkers would lead to the weaponization of AI in our threatcasting futures.



Simply because advances in science can be used to harm, does not mean that we should forgo their study. Rutherford's work also led to the discovery and use of cancer treatments, medical diagnostic tests still used today, and many other cases beyond military use. Today, machine learning has helped to improve our lives greatly, identifying spam email and preventing your computer from being infected, facilitating our ability to search anything, expanding knowledge to regions without proper education systems, early detecting of diseases through pattern recognition, and thousands of other use cases.

Funding basic research of machine learning and artificial intelligence should be a priority for the United States government. In 1941, the Manhattan Project and UK's Tube Alloys efforts were created and resourced by their respective governments to promote the science and achieve technological breakthroughs before Hitler could. Today, we once again face a technological race to achieve breakthroughs before our adversaries. We must return to funding basic and applied research in this domain to encourage research amongst academics and industry and to share technological advances in order to advance the body of knowledge in a positive way.

Government funding would promote an open source environment for the sharing of knowledge and understanding. While there is some open source within the tech industry many developmental breakthroughs and learnings are still proprietary. Corporations are allowed a tax credit for any dollar that they claim is spent on R&D, but they do not have to share the results. Academia loves to share research and ideas but is hampered by a lack of access to datasets and compute power. What if a company like Amazon could get tax credits for all the compute time that they share with academia for Artificial Intelligence / Machine Learning research instead of receiving tax credits for the money that they spend on their own R&D?

H.G. Wells wrote *The World Set Free* in 1914 - a futuristic novel about the destructive use of atomic weapons, however, this did not cause a stop to domain research. We should not allow today's media, storytellers, or even this report from delaying advances in artificial intelligence. The goal is to empower development of the potential benefits of this technology while remaining cognizant when making design decisions of potential threats as the technology evolves. ●

Appendix 1

THREATCASTING METHODOLOGY

While the threatcasting methodology was briefly discussed at the beginning of this report, this appendix provides more details to inform “how the sausage is made”.

The key to the process is the people. Participants come with a range of experiences, expertise, education, and passion. They are pre-assigned into 3-4 person groups for the duration of the process. The groups are specifically curated to take advantage of the diversity within the larger group. This small group assures that every member can express her/himself. Also the small group size allows for in-depth discussion and debate.

A fundamental component of the threatcasting process is selecting the appropriate research inputs to feed the future modeling. These focus themes are selected to explore how their evolution from today contributes to the future but also how the intersection of the focus areas’ growth modify each other. To select these themes, senior leaders inside the problem space and thought leaders outside the problem space are consulted on what “keeps them up at night” or what they feel no one is focused on yet to determine the severity and urgency of the proposed themes.

Next we curate and find SMEs to inform and bring these focus areas to life within the threatcasting sessions. These SMEs are individuals that can quickly describe the current state of their domain and how it might evolve over the next decade. They provide clarity to help participants hone and define threats in the future. Transcripts for the SMEs’ input are transcribed in Appendix 2.

THREATCASTING IS A FOUR PHASE METHODOLOGY.

Phase One: Research Synthesis

Research synthesis is the first phase of the threatcasting methodology. The purpose of this phase is to allow each small group to process the implications of the SME provided data while gathering the intelligence, expertise, and knowledge of the participants in the Research Synthesis Workbooks. These workbooks are located in Appendix 3.

During this phase, all participants listen to each SME’s presentation but they are assigned a specific presentation on which to take notes. At the conclusion of the presentations, they break into their assigned small group. Within these groups, they identify key elements and interesting points from their assigned presentation and conduct initial analysis. They explore, for each of

these points: 1) what the larger implication of that point would be within the future, 2) characterize this as either positive or negative, and 3) list ideas for what we should do about it. The “we” is purposely broad as the input can be personal to the small group, the collected team in the room, the entire company, or the entire human race.

The output of the research analysis phase is a numbered list of these key points from the SMEs as determined by participants.

Phase Two: Futurecasting

The core of the threatcasting methodology begins with phase two of the process. Each future is based upon the Research Synthesis Workbooks.

At the start of this phase, the participants return to their small groups and select a single data point from each of the SME presentations as described in the Research Synthesis Workbook roll-up. Groups make selections via random sampling with replacement for each SME. The instrument for sampling are 20-sided dice. Without this randomness, people often pick “easy” data points that fit with their view of the future. These points establish the framework of the future environment that they will model.

After establishing the visualization of the environment, the group imagines a specific person living in that future. The group envisions who the character is, whom their family is, and the broader community with which they identify. Then the group explores where the character lives, thinks about their occupation and visualizes what constitutes their normal way of life.

The physical or digital instantiation of the problem caused by the threat is the “event”. To better model and understand the event, the small group is asked a series of questions which are recorded in the worksheets in Appendix 3. Going beyond just the “5Ws” of traditional information gathering (who, what, when where, why) these prompts are specifically designed to create a more well-rounded narrative describing the threat.

Then our perspective changes and the groups see the event from the adversary’s perspective; exploring potential roadblocks or barriers and thinking about new business models and practices to enable the event. We imagine the technology that would help facilitate the threat and what support systems are required. Finally, we think about the training necessary to enable this threat. This change in perspective helps the small group to better define the threat, visualize the adversary’s motivations, and understand their desired end state that will be disrupted, mitigated and recovered from.

The end state of the futurecasting phase is that each small group has created a story about the future.

Phase Three: Backcasting

The third phase of threatcasting is the backcasting process. Here, still in these small groups, focused on the narrative they have created and the threat that they described – the groups think about what could be done to disrupt, mitigate, and/or recover from their defined threat actor.

During backcasting, there are two types of events that the groups explore. The first are gates. Gates are things that defenders (government, military, industry, etc) have control over that could disrupt, mitigate, and/or recover from the threat. These are things that will occur along the path from today to T+10 years. The second event type are flags. Flags are things the defenders don't have control over but once they occur, there is no going back. These flags should have a significant effect on the envisioned future. These are events we should be watching out for as heralds of the future to come.

Once the events are imagined, the small groups then timeline the actions to disrupt, mitigate, or recover from the threat. Thinking about the actionable objectives that need to occur in the next four years and also in the four years after that in order to protect against the future described threat. This iterative exercise gives the participants a chance to see how actions today can be built upon, achieving an interim goal and eventually guarding against the threat.

At the end of phase three, each small group reports out, telling the larger group a story about their person in a place with a problem. They describe the threat and what could be done to disrupt, mitigate and recover from that threat. Finally, the session ends with a discussion of the process and the collection of threats. The assembled group looks for patterns in the aggregated futures and also looks for areas that were not discussed. The session is concluded, leaving the entire group to continue to think about the futures.

Phase Four: Analysis and Final Report

Following the threatcasting session, the moderators use the Research Synthesis Workbooks as well as the small group Threatcasting Workbooks as raw data for a post-analysis. Reviewing each workbook, the team of moderators look for patterns in the futures and for areas that were not explored.

This synthesis exercise generates an aggregation of multiple futures and threats. Secondary research as well as the backcasting details from the practitioners give the team the raw data needed to make specific recommendations for near and long term actions to be taken. The final report collects the SME inputs, the participant worksheets and the team's post analysis. The post-analysis consists of multiple clustering and aggregation exercises to determine the patterns in all of the futures modeled during the event. These clusters are then examined in light of the SME presentations, looking for possible inconsistencies or areas that need more clarification. Additionally the team highlights areas that perhaps the groups did not model but were strong themes in the SME presentations. Combining all of these together, the team makes specific recommendations for next steps and areas of action, informed by the backcasting (gates, flags, milestones) provided by the participants.

Appendix 2

RESEARCH INPUTS

Six curated inputs from cross-industry experts helped inform the futures we modeled. First was Dr. Genevieve Bell, discussing how we should think about interrogating AI. Sam Harris posed the question of how we might build AI without losing control over it. Dr. Dave Gioe outlined 14 cyber considerations for humans. Paul Thomas discussed how to approach Threatcasting from an economic perspective. Andre LeBlanc outlined the growth, impact, and future of applying AI to real world industries. The sixth and final talk was MAJ Natalie Vanatta, PhD with a wrap up of key ideas from various expert interviews regarding cyber growth and our relationship with machines. Transcripts of all talks are located below.

The following research inputs were transcribed by machine and were not further edited. Some context might be missing or misplaced.

47

Dr. Genevieve Bell

Wow, I always wonder when I have to follow Tim O'Reilly about what on earth I'm gonna say, a man who quotes poetry and Luddites is a man after my own heart. I also realize, given what you've heard so far, I'm actually in a really nice place cuz I want to talk about the stories of AI, too, but from the point of view of an anthropologist. There's lots of ways of describing AI. I suspect you'll hear many of them on this stage over the next two days. For me, my favorite working definition at the moment comes from my colleagues, Kate Crawford and Meredith Whittaker, who defined AI recently in a publication as, "a constellation of technologies, including machine learning, perception, reasoning, and natural language processing."

You might reasonably ask, why would an anthropologist care about those things? Of course, the reality is AI is more than a constellation of technologies. You heard Tim talk about all the ways in which AI is woven into the human experience. I also want to suggest that it's a cultural category in and of itself. The fact that we can talk about the language over, "Is AI gonna take our jobs?" and we can refer back to the Luddite rebellion, tells us that AI is more than just a set of technologies. It's actually a cultural thing. It's a cultural category, and cultural categories, that's what anthropologists love. That's firmly in my wheelhouse.

Now the challenge here is, as an anthropologist, usually the way I like to make sense of things is I want to go hang out with them. We do fieldwork, right? We go spend time in the places where meaning is being made. We go spend time with people and ask them things. It's a little hard to think about how you would go do fieldwork with AI. I can think about doing fieldwork with the people who make AI. I can think about hanging out around algorithms and their impact, but I wanted to take a different tangent here and suggest that maybe what you could do instead was an ethnographic interview.

In anthropology, one of the ways of getting a sense of something is to conduct what we would call a semi-structured interview. [Audio cuts out 02:05] anthropologist named James Spradley who wrote the definitive book on doing this back in the 1970s. He said, if you really want to get to the basis of something, to the bottom of it, to how a person makes sense of things, you should ask them three kinds of questions. You should ask them descriptive questions to get them to talk about their world in their own language, you should ask them structural questions to get them to talk about how they make sense of their world, and you should ask them contrast questions so that you can work out what they think they aren't. Pretty simple, right?

For me, that meant there were five questions I wished I could ask AI. I'm gonna try and ask those questions and speculate what the answers might be. As a good anthropologist, the first question is: What's your name, and how did you get that name? Of course, asking that question of AI is interesting, right? The politics and the polemic, and even the etymology of the name itself, tell us something. I would argue it would only have been in the 1950s in America that artificial was a good thing. If we were naming AI now, we might not want to use the word "artificial." We consider that to be somewhat a problematic thing. We talk about organic and local and natural are the fetish objects, but if you think back to the 1950s, artificial was a good thing. It was human-made. It was different than the natural, which was wild and uncontrollable and messy.

Artificial had all the kind of shininess that Tim was just talking about of the postwar period. Artificial was about rubber and "al-u-min-ium," not aluminum, and about all those things, right? It was about a world we were making, and that was seen as being

good, and, of course, intelligence here is also about this notion of skills that can be acquired. It's about learning. It is, of course, always said in opposition to the emotional, the irrational; again, the messy. Here you have this really interesting contrast in the naming between two words that have an interesting relationship to each other. What does the difference between artificial and intelligence tell us, and what is going on in that naming convention? You could ask the same things about "machine" and "learning," and putting those two things together. The mesh of the human and the machine there is sort of a fascinating thing. That would be one kind of descriptive question.

A second one would be to go, "Who brought you up? Who raised you?" In the anthropological tradition, "Who are your mommies and your daddies?" In this case, it's a lot of daddies, I have to say, and, of course, the history of AI is equally complicated in its name. While many of us know "AI" was coined at a conference in 1956 at Dartmouth here in the United States, and many of the early founders of AI were at that conference, and they were mostly—well, in fact, nearly as I can tell, entirely men. They had very different preoccupations and concerns. They were radically interdisciplinary for the 1950s, had backgrounds in philosophy and mathematics, psychology, the emerging field of computer science.

While they had very different backgrounds and concerns, their notions of what it meant to be human and how humans learned was strongly influenced by behavioral psychology and by the behaviorists of the 1940s and 1950s, in particular a man named B.F. Skinner who had an idea about how humans worked that was about stimuli in and response out. That if you could track the stimuli, you could predict the

response, and if you changed the stimuli, you could change the response, what we would otherwise know as conditioning. Now, of course, if you want to think about building a machine, that is an excellent way to think about how humans work. Of course, Skinner was also in himself a moment reacting against Freud, the messiness of European psychology and the messiness of humanity. When you ask who raised artificial intelligence, it was raised by a very particular set of people with a very particular set of histories and funded by a very particular set of institutions and governmental agencies. Who raised AI? Good question.

You can also ask: Where did AI come from? Who are its people? Where's its country? As more of a structural question, Tim flagged some of those, right? AI has a long lineage, and many places it could call home before it was artificial intelligence. Human beings have been fascinated with making things come to life for thousands of years, I would say. We have myths inside most traditions about what it means to bring things to life, early human societies to the present day, lots of tales of that. They're mostly moral tales. They're mostly cautionary, and they frequently end badly. Think Gollum, Frankenstein, and the Terminator to just be one lineage. Of course, the practicalities of physical machinery were complicated, too, the Digesting Duck here, the Mechanical Turk, the Babbage machine. The question becomes: What are you mechanizing and why? What is the prize you are after about what makes something human or real or intelligent? Those are also questions not about just who raised you, but where you come from.

There is, of course, the question; again, a structural question of: What do you do every day? A good question I often ask when I'm doing fieldwork is: What's a typical day for you like? Well, we know a typical day for AI is a little complicated. We could argue if Hollywood told us AI is coming to kill us, the robot apocalypse and everything else. If we listen to the news, and Tim just gave you a lovely sense of the headlines, yeah, it's gonna replace us and our jobs. Of course, the reality is infinitely more complicated. We already live in a world riven through with algorithms, some of them benign and banal, from Amazon to eHarmony and Netflix, some of them more complicated around sentencing guideline tools and risk predictors.

We start to have the beginnings of semi-autonomous and autonomous machinery that are unevenly distributed and the regulations are complicated, and, oh, by the way, they all already have built into them a country and a culture. That there, to many engines, some of them Google's visual ones, reads as graffiti. If you're from Australia, you know that's a wall of a shearing shed, and those are the names of all the men who sheared there over the last hundred years. I can read that wall and tell you when the good seasons were and when the bad seasons were and when the war happened, but that's because I recognize that as a cultural object. As we start to think about how things are made sense of, culture here will turn out to matter, too.

Last, but by no means least, there are Spradley's classic contrast questions of: What aren't you? Can we ask an AI object what it dreams of? Could we talk about artificial subconscious or an artificial id? Would that even make sense? This object here is a painting rendered by a performance artist in Japan who hacked a Roomba and turned its dust-sucking vents into paint-blowing vents and imagined what the art would look like. When he did that, he's starting to speculate about: What would the artistic intent of an object be?

Back in the early days of AI, one of the founding participants, a man named McCarthy, actually argued that machines might have beliefs, and if they did, how would we manage them? He said, if they were sentient, they must also have, in some ways, a core set of principles which were remarkably akin to beliefs. If we ask a set of questions here and start to assume that AI isn't just intelligence, but beliefs or a subconscious, we should then also ask questions about what it might mean to have dreams and what it might

mean to have intentionality beyond the things that it is trained to do. Where does all of that leave us?

Well, for me, I think it starts by saying we ought to critically interrogate AI, not just as a technology, but as a culture, and as we make sense that there are a series of things that open up a broader set of conversations, if we understand the politics of its name, we can understand partly how it is that we think about it, and what it might mean to rename it and imagine it differently. If we know who its founders are, we can ask questions about what were there biases and their preoccupations, and what, as a result, might we need to bring back into the conversation, including, I would suggest, an approach to interdisciplinarity. If we understood the intellectual and cultural genealogies of the people and of the objects themselves, I suspect that opens up ways of managing fear, but, also, of new possibilities. If we think about the work of AI as also the work of humans, and of us making AI, that means we're encoding it with our own biases and normative thinking, and, frankly, as a good feminist, I wonder about a little bit more feminist theory, and perhaps a bit of queer theory here to unpack what it means to be normal.

Last, but by no means least, perhaps, when we talk about AI, we also ought to talk about things like love and fear and hate and our humanity, and perhaps even our souls, because as we do that, what becomes clearer is that AI is just another manifestation of what it means to be human, and putting our full human selves back into that conversation turns out to be critical. With that, I want to end and say thank you.

Sam Harris

Ted Talk

Can We Build AI Without Losing Control Over it?

I'm going to talk about a failure of intuition that many of us suffer from. It's really a failure to detect a certain kind of danger. I'm going to describe a scenario that I think is both terrifying and likely to occur, and that's not a good combination, as it turns out. And yet rather than be scared, most of you will feel that what I'm talking about is kind of cool.

I'm going to describe how the gains we make in artificial intelligence could ultimately destroy us. And in fact, I think it's very difficult to see how they won't destroy us or inspire us to destroy ourselves. And yet if you're anything like me, you'll find that it's fun to think about these things. And that response is part of the problem. OK? That response should worry you. And if I were to convince you in this talk that we were likely to suffer a global famine, either because of climate change or some other catastrophe, and that your grandchildren, or their grandchildren, are very likely to live like this, you wouldn't think, "Interesting. I like this TED Talk."

Famine isn't fun. Death by science fiction, on the other hand, is fun, and one of the things that worries me most about the development of AI at this point is that we seem unable to marshal an appropriate emotional response to the dangers that lie ahead. I am unable to marshal this response, and I'm giving this talk.

It's as though we stand before two doors. Behind door number one, we stop making progress in building intelligent machines. Our computer hardware and software just stops getting better for some reason. Now take a moment to consider why this might happen. I mean, given how valuable intelligence and automation are, we will continue to improve our technology if we are at all able to. What could stop us from doing this? A full-scale nuclear war? A global pandemic? An asteroid impact? Justin Bieber becoming president of the United States?

The point is, something would have to destroy civilization as we know it. You have to imagine how bad it would have to

be to prevent us from making improvements in our technology permanently, generation after generation. Almost by definition, this is the worst thing that's ever happened in human history.

So the only alternative, and this is what lies behind door number two, is that we continue to improve our intelligent machines year after year after year. At a certain point, we will build machines that are smarter than we are, and once we have machines that are smarter than we are, they will begin to improve themselves. And then we risk what the mathematician IJ Good called an "intelligence explosion," that the process could get away from us.

Now, this is often caricatured, as I have here, as a fear that armies of malicious robots will attack us. But that isn't the most likely scenario. It's not that our machines will become spontaneously malevolent. The concern is really that we will build machines that are so much more competent than we are that the slightest divergence between their goals and our own could destroy us.

Just think about how we relate to ants. We don't hate them. We don't go out of our way to harm them. In fact, sometimes we take pains not to harm them. We step over them on the sidewalk. But whenever their presence seriously conflicts with one of our goals, let's say when constructing a building like this one, we annihilate them without a qualm. The concern is that we will one day build machines that, whether they're conscious or not, could treat us with similar disregard.

Now, I suspect this seems far-fetched to many of you. I bet there are those of you who doubt that superintelligent AI is possible, much less inevitable. But then you must find something wrong with one of the following assumptions. And there are only three of them.

Intelligence is a matter of information processing in physical systems. Actually, this is a little bit more than an assumption. We have already built narrow intelligence into our machines, and many of these machines perform at a level of superhuman intelligence already. And we know that mere matter can give rise to what is called "general intelligence," an ability to think flexibly across multiple domains, because our brains have managed it. Right? I mean, there's just atoms in here, and as long as we continue to build systems of atoms that display more and more intelligent behavior, we will eventually, unless we are interrupted, we will eventually build general intelligence into our machines.

It's crucial to realize that the rate of progress doesn't matter, because any progress is enough to get us into the end zone. We don't need Moore's law to continue. We don't need exponential progress. We just need to keep going.

The second assumption is that we will keep going. We will continue to improve our intelligent machines. And given the value of intelligence -- I mean, intelligence is either the source of everything we value or we need it to safeguard everything we value. It is our most valuable resource. So we want to do this. We have problems that we desperately need to solve. We want to cure diseases like Alzheimer's and cancer. We want to understand economic systems. We want to improve our climate science. So we will do this, if we can. The train is already out of the station, and there's no brake to pull.

Finally, we don't stand on a peak of intelligence, or anywhere near it, likely. And this really is the crucial insight. This is what makes our situation so precarious, and this is what makes our intuitions about risk so unreliable.

Now, just consider the smartest person who has ever lived. On almost everyone's shortlist here is John von Neumann. I mean, the impression that von Neumann made on the people around him, and this included the greatest mathematicians and physicists of his time, is fairly

well-documented. If only half the stories about him are half true, there's no question he's one of the smartest people who has ever lived. So consider the spectrum of intelligence. Here we have John von Neumann. And then we have you and me. And then we have a chicken.

Sorry, a chicken.

There's no reason for me to make this talk more depressing than it needs to be.

It seems overwhelmingly likely, however, that the spectrum of intelligence extends much further than we currently conceive, and if we build machines that are more intelligent than we are, they will very likely explore this spectrum in ways that we can't imagine, and exceed us in ways that we can't imagine.

And it's important to recognize that this is true by virtue of speed alone. Right? So imagine if we just built a superintelligent AI that was no smarter than your average team of researchers at Stanford or MIT. Well, electronic circuits function about a million times faster than biochemical ones, so this machine should think about a million times faster than the minds that built it. So you set it running for a week, and it will perform 20,000 years of human-level intellectual work, week after week after week. How could we even understand, much less constrain, a mind making this sort of progress?

The other thing that's worrying, frankly, is that, imagine the best case scenario. So imagine we hit upon a design of superintelligent AI that has no safety concerns. We have the perfect design the first time around. It's as though we've been handed an oracle that behaves exactly as intended. Well, this machine would be the perfect labor-saving device. It can design

the machine that can build the machine that can do any physical work, powered by sunlight, more or less for the cost of raw materials. So we're talking about the end of human drudgery. We're also talking about the end of most intellectual work.

So what would apes like ourselves do in this circumstance? Well, we'd be free to play Frisbee and give each other massages. Add some LSD and some questionable wardrobe choices, and the whole world could be like Burning Man.

Now, that might sound pretty good, but ask yourself what would happen under our current economic and political order? It seems likely that we would witness a level of wealth inequality and unemployment that we have never seen before. Absent a willingness to immediately put this new wealth to the service of all humanity, a few trillionaires could grace the covers of our business magazines while the rest of the world would be free to starve.

And what would the Russians or the Chinese do if they heard that some company in Silicon Valley was about to deploy a superintelligent AI? This machine would be capable of waging war, whether terrestrial or cyber, with unprecedented power. This is a winner-take-all scenario. To be six months ahead of the competition here is to be 500,000 years ahead, at a minimum. So it seems that even mere rumors of this kind of breakthrough could cause our species to go berserk.

Now, one of the most frightening things, in my view, at this moment, are the kinds of things that AI researchers say when they want to be reassuring. And the most common reason we're told not to worry is time. This is all a long way off, don't you know. This is probably 50 or 100 years away. One researcher has said, "Worrying about AI safety is like worrying about overpopulation on Mars." This is the Silicon Valley version of "don't worry your pretty little head about it."

No one seems to notice that referencing the time horizon is a total non sequitur. If intelligence is just a matter of information processing, and we continue to improve our machines, we will produce some form of superintelligence. And we have no idea how long it will take us to create the conditions to do that safely.

Let me say that again. We have no idea how long it will take us to create the conditions to do that safely.

And if you haven't noticed, 50 years is not what it used to be. This is 50 years in months. This is how long we've had the iPhone. This is how long "The Simpsons" has been on television. Fifty years is not that much time to meet one of the greatest challenges our species will ever face. Once again, we seem to be failing to have an appropriate emotional response to what we have every reason to believe is coming.

The computer scientist Stuart Russell has a nice analogy here. He said, imagine that we received a message from an alien civilization, which read: "People of Earth, we will arrive on your planet in 50 years. Get ready." And now we're just counting down the months until the mothership lands? We would feel a little more urgency than we do.

Another reason we're told not to worry is that these machines can't help but share our values because they will be literally extensions of ourselves. They'll be grafted onto our brains, and we'll essentially become their limbic systems. Now take a moment to consider that the safest and only prudent path forward, recommended, is to implant this technology directly into our brains. Now, this may in fact be the safest and only prudent path forward, but usually one's safety concerns about a technology have to be pretty much worked out before you stick it inside your head.

The deeper problem is that building superintelligent AI on its own seems likely to be easier than building superintelligent AI and having the completed neuroscience that allows us to seamlessly integrate our minds with it.

And given that the companies and governments doing this work are likely to perceive themselves as being in a race against all others, given that to win this race is to win the world, provided you don't destroy it in the next moment, then it seems likely that whatever is easier to do will get done first.

Now, unfortunately, I don't have a solution to this problem, apart from recommending that more of us think about it. I think we need something like a Manhattan Project on the topic of artificial intelligence. Not to build it, because I think we'll inevitably do that, but to understand how to avoid an arms race and to build it in a way that is aligned with our interests. When you're talking about superintelligent AI that can make changes to itself, it seems that we only have one chance to get the initial conditions right, and even then we will need to absorb the economic and political consequences of getting them right.

But the moment we admit that information processing is the source of intelligence, that some appropriate computational system is what the basis of intelligence is, and we admit that we will improve these systems continuously, and we admit that the horizon of cognition very likely far exceeds what we currently know, then we have to admit that we are in the process of building some sort of god. Now would be a good time to make sure it's a god we can live with.

Thank you very much.

.....

Dr. David Gio

Hi. I'm Dr. David Gio. I'm the history fellow at the Army Cyber Institute at West Point, but today I speak to you as a former CIA operations office and also as a former chief of intelligence and counterintelligence for a deployed army combined joint task force.

Today I'd like to spend a couple minutes talking to you about the impact of cyber considerations on human intelligence. I've got 14 points in honor of Woodrow Wilson. I know that you're rolling your eyes right now. All I can say is that in 1919 when

Wilson brought his 14 points to Paris, the French Prime Minister George Clémenceau said, “Woodrow, even the Good Lord himself only had 10 points,” but I’m gonna stick with my 14 in honor of Woodrow Wilson. I’ll move through them expeditiously.

Point 1. Cyber considerations matter in human intelligence, which I’ll HUMINT for this brief presentation. We can’t just stick our heads in the sand and hope that they’ll go away or that the cyber thing will blow over. Given the pervasive nature of the digital global terrain and the interconnected and networked world, cyber considerations must at least play a part of any successful human operation.

Point 2. HUMINT never went the way of the dodo. The idea that the National Security Agency became so good and so capable in cyberspace that CIA simply gave up on traditional HUMINT recruitment in the digital age is nonsense. The traditional HUMINT recruitment cycle of spotting or identifying, assessing, developing, recruiting and then handling agents still pertains even today in the cyber age and HUMINT remains effective against all targets.

Point 3. Some of the HUMINT recruitment cycle can be done online. It’s amazing what people will put online, myself included. For instance, spotting new sources can now be done via online trolling. People today volunteer loads of information about themselves and, under the guise of professional networking, often reveal much about their so-called placement and access, that is to say the position that they hold and what they have access to on a routine basis. This can save much time by shortening or truncating some aspects of the HUMINT recruitment cycle, but it doesn’t remove them entirely and you still have to go through the whole thing.

Point 4. Cyber considerations have changed the character but not the nature of HUMINT. What do I mean by that? There’s simply no substitute for personal interaction and the kind of relationship that a case officer and his or her agent develop can’t be done solely online. Even if the reporting of information itself could be done virtually, HUMINT needs rapport and trust, and that’s harder to achieve across a network.

Point 5. There’s no doubt that cyber considerations will impact working undercover and counterintelligence precautions to preserve cover will have to withstand a level of hostile scrutiny that would’ve been unimaginable 20 years ago. It will be increasingly challenging to maintain a cover legend where digital detritus is ubiquitous or maybe should be, and almost every aspect of one’s cover legend can be checked almost instantaneously by anyone with a Google search. If you were challenged as to what does your commute look like in the morning and you were saying that you live in a city that you’ve never actually visited, that’s gonna be a problem when, with a combination of Google Earth and Yelp, you could figure out where that person might go to lunch on any given day. If they don’t know where that is, that’s gonna be a problem for your cover persona. A thought question along these lines for the threat cast is: Should you develop a Facebook or a LinkedIn account through your cover persona or does that create more problems than it solves?

Point 6. Big data matters for intelligence analysis and harnessing its power should be a key task of any serious intelligence community. Open source intelligence, also known as OSINT, and social media intelligence, a new phenomena now known as SOCMINT, are going to matter as collection mechanisms more than in the future than they do now, but this will only take them to co-equal status with HUMINT and SIGINT instead of outperforming the more traditional intelligence types.

Point 7. Like most emerging technologies, biometrics will be a double-edged sword. The military loves biometrics because they can use biometric means to identify, record, and track enemy combatants or suspected enemy

combatants, particularly in counterinsurgency type environments such as Afghanistan and Iraq. CIA operations officers on the other hand hate biometrics because it poses significant operational challenges to working an alias and particularly crossing borders. A thought question on this point might be: Should CIA try to defeat biometrics via other uses of technology or simply move along to more permissive operational environments?

Point 8. Communication between a case officer and his or her agent has always been the long pole in the tent, and historically most technological developments in HUMINT have come in the communications arena. Cyber means have eased and enabled agent communication in profound ways. Commercially-encrypted online communications are faster than ever and offer unprecedented levels of security.

A corollary to that is Point 9. Case officers can't lose sight of the human in HUMINT. Although cyber advancements have eased and facilitated agent communication, online or virtual communication, no matter how frequent, is not a substitute for person-to-person contact. There's no algorithm or online chatroom that can replace a case officer's gut feelings about an agent's reliability, veracity, or motivations, and the agent might find that he feels at the very far end of a fiber optic cable without that personal interaction.

Point 10. There's an old saying in HUMINT that if something is convenient it probably means you're doing it wrong or it's operationally insecure. Cyber tools are powerful indeed, but the line between convenience and shortcuts is sometimes hard to identify. Therefore, operational managers will have to pay attention to ensure subordinates aren't taking virtual

shortcuts in handling their human assets.

Point 11. Anyone with a security clearance is vulnerable to a hostile intelligence service getting loads of personal data about that person. It doesn't matter how many precautions you take to carefully manage your personal privacy; there are too many weak points beyond your control. For instance, take the 2015 hack of the U.S. Office of Personnel Management, which resulted in the theft of hundreds of thousands of federal employees SF 86 security clearance forms, including my own, that had personal and private data.

The corollary to Point 11 is Point 12. Just because a hostile intelligence service knows a lot about you through a massive hack or a personalized and tailored spearfishing campaign doesn't mean ipso facto that they will be able to recruit you or pressure you to reveal classified information. They would need to weaponize and exploit that information, which isn't as easy as it sounds.

Point 13. Massive leaks of classified information such as those done by Edward Snowden and Chelsea Manning are absolute kryptonite to HUMINT operations. Leaks erode trust between a case officer and his agent who is promised that his or her information would be well protected at all costs. In an era of WikiLeaks and similar mass leaking platforms, the job of the case officer is exponentially harder because he will have little credibility about source protection.

Point 14. In sum, cyber tools will not replace or supplant traditional HUMINT operations, but it will make some aspects harder, some aspects easier, and will present some new challenges for consideration. The successful case officer of the future will be the one who best understand these new capabilities and limitations, and the intelligence winners will be those who adapt the fastest to a rapidly-changing world.

Thanks a lot for your attention and enjoy your threat casting workshop.

.....

Paul Thomas

Please Note: There were some audio problems with the source video. The inaudible version has not been transcribed.

Interviewee: Hi. My name's Paul Thomas. Ryan and I work together at Intel. I was chief economist for Intel Corporation until I retired in June last year. I think the work that Ryan's been doing sounds pretty exciting and he asked if I could help by just making a few comments. He hasn't told me a lot about the [inaudible 00:30].

My first comment is that by concentrating on threat testing, [inaudible 00:42] term, you're almost necessarily focusing on what you're supposed to which is threats, damage to our economy, to our national defense, to our way of life, malevolent actors, a lot of mistakes and threats may occur through accidents and we're looking at that option.

My first comment is a threat testing by focusing on the negative does create possibility that you won't pay enough attention to the positives. One way to avoid risk is to avoid exposures that could generate [inaudible 01:30], and so just keep in mind that there are threat testing, opportunity testing, picture testing. For example, people who pay taxes talk as if they're [inaudible 01:46] investments take so much [inaudible 01:48] in practice if they're taking good care of the portfolio. They may let the importance of taxes affect the behavior but they should [inaudible 02:01] games. There are examples in [inaudible 02:07] economy where risk law can [inaudible 02:09] so you may not look at those as [inaudible 02:12] discuss those possible loss opportunities at some point, I think.

The second point is you're doing inclusive modeling. You're probably building interior models [inaudible 02:21] supposed to leave paper or you're just thinking in those terms. I would say [inaudible 02:23] it's okay to start with some of those. Just simple model [inaudible 02:36] interesting behavior in the face of risk and uncertainty. With simple models, [inaudible 02:39] you don't have a lot of human opponents or players within the formed coalitions or who could see threats. You instead try to figure out

how to get through when you're--or how to get through tsunamis and earthquakes and some accident [inaudible 03:01] technology. You can do that and you can understand some of the ways that people [inaudible 03:09] mitigate them or in some cases voluntarily expose themselves to some risks [inaudible 03:14] you can see this in insurance [inaudible 03:17].

In the insurance literature there's a given called the Bernoulli Principle. Bernoulli [inaudible 03:25] in so many things. Physics [inaudible 03:30] but also statistics [inaudible 03:29]. The Bernoulli Principle when it comes to insurance says if you can get fairly priced [inaudible 03:36] fairly priced insurance, do it. If you can sort of keep the return to more or less constant to avoid risks. If you're risk averse, you'll give up some gains to avoid risk even more.

There are reasons that anybody in society should not act as [inaudible 03:56]. We already have pulled a lot of our resources in the country, the nation, other nations and coalitions were [inaudible 4:05] again and form so much of the risk in certain literature. Part of what he found with what they did with the land is that you don't necessarily want to do the things that an individual would consider wise as a country who have insurance [inaudible 02:26]. There are other lessons that will come from that principle [inaudible 04:31].

The second point I wanna make is that once you start having gains that can set the players you open yourself up to interesting behavior coalitions. Probably once you think about our allies and adversaries around the world. You could see that in some ways, we're all rivals. In some ways we cooperate. We're not [inaudible 05:07] there are some gains. Sometimes we have some gains from working together and

cooperating. Sometimes, we've had some gains from networking [inaudible 05:17] the agreement. We certainly disagree on many trade issues with some of our closest nations [inaudible 05:27].

The effect you get when you actually have multiple players, complexity. Another thing we get is the value of being able to bind yourself. US is usually the biggest, most powerful player in any particular phase of the continuing game we play but other countries might worry that we will motivate some sort of big giant [inaudible 06:03] question. It absolutely will be a constraint [inaudible 06:07] behavior if put enough online [inaudible 06:12] hurt us, she can trust us. This happens in companies in the corporate world for one of the things that [inaudible 06:21] electrical utilities that [inaudible 06:24] coal. We'll put their power plants near the coal mine openings and there's two reasons for that. One, it's cheaper to transmit a chain of coal.

Secondly, you don't have to worry that if you specialize your output for a giant utility customer that they're going to say, "We're not going to buy from you. You have to lower your price. It's [inaudible 06:46]." You can say, "Hey you put your utility plant right by my mine mouth, you're vulnerable too, and that starts off concerning [inaudible 07:00]. That's strange when you think about risk, you actually expose yourself to risk to buy credibility with your friends and partners.

The last thing that I'll say and I wanna get this as short as I can is that markets express themselves in various ways and they can be embedded in multiplayer games and the work has been done. Doesn't always give additional insight but it tells you that if you're thinking

about games, you better think of the markets and I just wanted to point out that there are risks you probably [inaudible 07:37]. An example would be the risk of [inaudible 07:41] water and [inaudible 07:43] expect a lot of suffering and perhaps protests and disorder and [inaudible 07:49] movements. You may be worried about the fact that China dominates the production of railroads and that they could use some of those elements in their weapons and electronics. That's overexposed.

You're probably worried about changes. Markets are pretty good [inaudible 08:13] so sometimes the recipe would be [inaudible 08:18]. As an example, we're about to run out of oil for the last 100 years with people saying mostly we have 40 years of oil and we always seem to have 40 years of oil. Using current definitions and so forth, you could preferably extract [inaudible 08:34] for the introduction of more technology, we'll probably always 400 years of oil and the same things gonna be true about natural gas. I'm not trying to argue therefore you should try alternative fuels. I'm trying to argue that we might run out of tin, we're gonna run out of copper, we're gonna run out of coal, we're gonna run out of steel. Various times we'll [inaudible 08:57] risks, but realize the markets which crop up all the time there. An experiment in economics, it's very hard to keep markets from appearing in the experimental settings that they've seen up here. Keep that in mind.

Please let Ryan know if you have any questions for me and perhaps we'll have a chance to meet in [inaudible 09:21]. Thank you, bye.

.....

Andre LeBlanc:

https://www.youtube.com/watch?v=xH_B5xh42xc

Speaker: I'm going to talk to you guys about artificial intelligence, and artificial intelligence is still—I'm very excited about it. I'm very passionate about it. It's growing at an exponential rate, and what's exciting about it is we're not even at one percent of what artificial intelligence is gonna be, but it is an exponential trend, so by the year 2035, a lot of experts are saying that computers will be just as intelligent as a human. That's in about 20 years. By 2045, Ray Kurzweil has projected that computers will be more intelligent than all of humans combined. This is an exponential trend.

Just to give you a little bit of background about the evolution of technology and everything, we used to develop technology to replace our muscles, right? To do it's pick up a bigger boulder, a bigger truck, a bigger this, a bigger that. A lot of physical things have been replaced by technology in the past. In the last hundred years, we've been replacing—we've been doing things to enhance our brains. The calculator would be a simple example of enhancing our ways of doing math. Then we started automating things that were repetitive. Google's another type of artificial intelligence in the sense that if I had Google, and nobody else here did, I would seem like a very intelligent person. What this does is it enhances our experience, but it's not direct artificial intelligence.

Now, if you look right now at the stock market, there's one company in the US, they do 500 million automated trades per month. This is a robot that takes tons of information, makes decisions every millisecond, and makes hundreds of millions of dollars every year. This is just an example of some of the artificial intelligence that's coming. This is gonna happen in every industry. There's no way that a human could do that. There's no way that a human could compete with that, because there's so much information going through the system, and it's moving at the speed of light.

The next generation of AI is going to be more adaptive. It's going to be self-learning and it's going to be intuitive. When things

change, that's when automation kind of fails. The next generation of automation, what's gonna happen is, if something changes, it'll be able to change its own rules. Just as an example, I have an artificial intelligence company and what it does is, if I give it a simple command—like if I told a person to log into a website, intuitively we know how to log into a website. We know what to do. It's the same thing. If you can record that pattern of behavior, once you've seen enough websites, a computer can do that as well. Eventually, basically anything that's done on a computer will be able to be emulated. Eventually, these computers will be so intelligent that it will lead to the singularity, which is what Ray Kurzweil calls it. That will mean that the human race as we know it will become obsolete. Exciting, isn't it?

A lot of people think this is what's gonna happen when this happens, but the reality is this has happened before. Two-hundred years ago, 90 percent of people worked in the agricultural age. They worked in agriculture. Now two percent of people work in agriculture. Now, are we better off or are we worse off? We're better off. Things got better. The same thing's gonna happen. In the next 40 years as all these technologies get faster and faster, we're gonna have to shift to do something else, but it's gonna be good for the human race in general, because the robots will be working for us.

What's also important to understand is the AI world will be virtual. Most people see AI as a robot, but really Google—when you do a Google search, millions of algorithms are running in the background on servers somewhere else. It's gonna be the same idea, but AI are gonna be doing tests in a virtual world. We can call that the matrix. We can do whatever we want, but what's gonna happen is that if you wanted

to find the cure for cancer, what more effective way to do it than to test on a simulated human being a billion times with a certain drug, right? Instead of doing it on a rat, or on a monkey, or on all these things, you're gonna do it in a simulated environment a million times. Let's say this AI is look for a cancer drug, this one's looking for a Parkinson drug. This AI will develop a theory on trying to find a cure for that disease, and it may find something that helps out the Parkinsons. These billions of AIs will all be working together, and in the next generation, most inventions and most cures in medical fields will be found by AI and not by humans.

Now, when's this gonna happen? Well, Bill Gates said, "We always overestimate the change that will occur in the next two years, and we underestimate the change that will occur in the next ten." The reason he says that is because of exponential growth. If you think about Moore's Law, where technology doubles every two years—fifty years ago, if you had a cellphone—if you have a cellphone now, fifty years ago it's more powerful than all the computers combined. Computers are getting faster and faster and faster, and they're getting cheaper and cheaper and cheaper, and so 25 years from now, they're projecting that computers will be able to—a cellphone-sized computer will be able to fit into a cell, into a blood cell, so you'll be able to take a vaccine and maybe it'll inside and repair your body. Who knows? Nanotechnology's growing at an exponential rate as well. Nanotechnology, medical technology, all these technologies are growing exponentially and we're not even in the baby phase of what's going on. These trends are gonna to start to go up, and things are gonna change drastically and very quickly.

Fifteen-hundred years ago, in the past, before the year zero, it took about 1,500 years to double knowledge. Every 1,500 years you would double knowledge, so you'd have two times the knowledge, and then every 1,500 years after that, double again. We started writing books, and we started to write things down in language and things like that, and it started to happen every 250 years. Then we discovered science, and then it started happening every 25 years. Then we created information technology, it happened every eight years. Now we're at the information age and it's happening every year. Literally every 12 months, human knowledge is doubling. We're getting to a point

where the human brain can't comprehend all the information that's coming in. What's gonna happen is—200 years ago, you could have been an expert in ten different fields if you studied them enough. Now you can barely be an expert in one field. You have to go into a section of a section of a field and be an expert in that, and to make an impact in that field will be very difficult.

The human brain has about 86 billion neurons. It's about four times as much as a chimpanzee. It's not enough. This is where artificial intelligence is gonna come in. All this information that's coming in is gonna get managed by artificial intelligence, but at the same time the human kind is gonna learn from this AI. What if we could add ten times the neurons that you have in your brain? What if you could add a hundred? What if you add a million, a million times the neurons that you have right now, into the cloud, wirelessly connecting? Imagine if you had a doctor that says, nah, I don't like technology, the internet's not for me, and then you had another one where information is doubling every year—in fact, it's gonna double every 12 hours at one point, according to IBM. Imagine one month goes by. That one month has doubled sixty times, so you'd need to take an eight-year degree to learn what happened last month. Right? Not gonna happen. So the doctor that's connected to AI, he's gonna walk into the doctor's office, he's gonna run a windows update. Boom. He's learned everything that's happened in the last month, and this is how things are gonna happen. They're gonna accelerate at a dizzying pace and it's gonna be incredible.

My last slide here is The Matrix. At some point we're gonna have to make a choice. You take the blue pill and you believe whatever you wanna

believe. I don't believe that people will be forced into this system. You know, if you don't want to live in reality, you could go and play something. It already exists. It's called videogames, right? If you wanted to be in World of Warcraft and be a fantasy fighter and do all that stuff, you could do it in a virtual world, fill your boots, but if you want to take part in the next 50, 100 years, which—forty years from now, we will look back to today and say, I cannot believe we used to do that. In the medical field, in every field possible, we'll look back and we'll say, I cannot believe the amount of advances that happened. In the next 30 years, there's gonna be just as much change happened that happened in the last 2,000. Technology is exploding, but it's gonna be good. It's gonna be good for all of us. We're actually gonna benefit from it. We might even be able to work less, right? We used to work 80 hours a week when we used to work agriculture age. Now we work 40. Maybe we'll work 20, right? I probably still work 100, but that's just me. It's gonna be an exciting time and I hope you're excited about the future. I'm excited about the future. Thank you very much.

Dr. Natalie Vanatta

Good morning. I have a few points to add to your threatcasting design today as you think about it. A couple things to think about as you're imagining the visions of the future and what life might actually look like in 2027.

I'm gonna start with Moore's law. Just I feel we haven't talked enough about it yet. Which is just merely the observation that the number of transistors in a dense integrated circuit approximately doubles every two years. Okay, that's cool. Now, many people will argue that Moore's law is eventually going to die. It can't continue to work. Yet, Intel's latest chip set that they released this—earlier this year has transistors that are 10 nanometers wide. Truly amazing. Really, truly, it's not that law—Moore's law is gonna fail. Moore's law is gonna fail because of physics. Our understanding of physics today is what is gonna cause Moore's law to fail.

We still anticipate that by 2027, in the future that you're gonna be describing, that we'll see transistors on chips that are five

nanometers wide. To put it in an easier context, that's 12 atoms. Twelve atoms is how wide transistors will be on this—on chips in the future that you're imagining. Which means computers are gonna be everywhere, and everything can be a computer. If we add that in with a world where everything is censored, because we've transitioned from this idea of Internet of Things to internet of everything, how does the world completely change? In fact, are we not living in a computer at that point? It's truly only gonna be our imagination that's gonna limit what we're gonna be able to do at that point.

Now consider the fact that your phone—I mean, mine's a crappy phone, so definitely your phone. Your phone today—sitting in your pocket, hopefully, right, 'cause you're not checking it, 'cause you're paying attention to me—has more computational power than what we used to send a man to the moon. That's what you have on you right now, what we were able to do then. I challenge you to think about what's next. In 2027, as you're looking at your person in a place doing a thing, what are we using this power to do that we're surrounded by? Think about, what was science fiction even 10 years ago, which is gonna be reality in another 10? How does what we're using this immense power for influence or shape your person in a place and what they're doing and, more importantly, how they're actually living?

If we think about we're in this world where sensors are everywhere and we're surrounded by compute power, let's think a moment about how operations or businesses are actually conducted. Whether you're looking at it from a military perspective or a business perspective, or even from the perspective of the threat that's gonna be attacking your future vision, I would

say that there are generally three different ways that we conduct operations. I'm gonna call them normal, covert, and clandestine. Just so that we're all on the same sheet of music with regards to terminology, in a covert operation that I'm gonna conduct, it's planned and executed as to conceal the identity or permit plausible deniability by the sponsor. What's that mean? The bad guy doesn't want anyone to know that it was them or who sponsored the event. It's okay if the world knows it happened. In a covert operation, we don't want anyone to know who actually did it.

Completely different from the idea of doing operations in a clandestine mode, where the operation is sponsored or conducted in such a way to ensure secrecy, which means we don't want anyone to know about it at all ever. Now whether we're—that target audience that we don't want to know is the general public or some very specific entities, in a clandestine mode, you wouldn't even want to know that the operation occurred. Which then leaves everything else to be normal mode, right? That sets up some interesting things, I think, to ponder. In a world we're living in, where everything is monitorable and we're surrounded by compute power, can anyone truly ever conduct covert or clandestine operations ever again? Can you keep things a secret at all anymore?

Whether this is a business that's trying to get the skinny on their competitors or a nation state trying to do something against an adversary, think about it as you envision the threat. Were there actions that they did, the operation that your threat conducted against your person in a place doing a thing, where they intended to be covert or clandestine? If they were one or the other, was that because they were forced to be one or the other, 'cause you can't do one? Is it because the others were difficult? Is it a deliberate choice for how we were operating it? More importantly, does it matter anymore in your vision in the future if we can or cannot keep secrets anymore?

I want to pull that thread just a tad bit more. Not just the fact of how do you conduct operations, but think about the preparation activities that your threat is going to take prior to the start of the operation. Whether you're thinking about it from a military perspective, some military would call this intelligence prep

for the battlefield, or you're thinking about it from a technical perspective—pick any hacking methodology out there, how is your threat gonna scan the space and enumerate the weaknesses? What sorts of data or information will be present in this environment to enable these preparatory activities to occur? How is your threat going to have access to this information? Can they access this in a covert way? Can they access this in a clandestine way? Is it just open and available for all? Does it matter anymore if everybody sees what everyone's doing? Ultimately, how will this collection of information be turned into actionable intelligence?

That leads to a fourth point to bring out. In general, we've talked about it a little bit today. We normally start with data. Data is just the facts of the world. Then we process that data. Today we might do it by human or by machine. This becomes the information. Then we take and we analyze and refine this information, either by human or machine today. Now we have knowledge. Today, after we have knowledge, we use human judgment to turn that knowledge into situational understanding for us to understand the why. By having situational understanding, now this enables leaders to make decisions.

The question is that's today. What's going on in 2027 in your world? In everything that the experts have talked about this morning already, we're gonna be in a world where we're gonna need to be able to make decisions at machine speed, and AI will be prevalent. How is this gonna happen? How have we, between now and 2027, attempted to increase the cognitive capacity of our leaders and our workforce to be able to handle all this information and handle this notion of compute power so that they

can make decisions quicker? No matter about your personal belief when or if machines are gonna take over the world, I can guarantee it's not gonna happen in the next 10 years. We're only looking 10 years out. What does our relationship look like with them, with artificial intelligence or robots?

Have we made a determination that there's a certain subset of tasks that humans are just really good at doing? Are there a certain subset of tasks that machines are just really good at doing? How have we shaped that conversation? How have we continued the evolution of the human—of humankind in our learning to be better postured in this space, to let the humans do what humans do best and let machines do the rest? What does this look like in your vision of the future for a person in a place doing a thing?

Malcolm Gladwell's theory today is that it takes 10,000 hours of practice to achieve mastery in a field. In the future, have we figured out a way to either compress this expertise or this experience so that younger leaders are empowered to make decisions? Have we thought about how we're gonna optimize human performance in this new future set that we're in? Have we thought about what are those skills that are necessary to be successful in the society of computers and information?

My final thought for you this morning is the question of, do we potentially need programmer archaeologists in 2027? I use that term very specifically 'cause it's absolutely not mine. It's from an awesome piece of science fiction entitled *The Darkness in the Sky*. The notion in that book is that the future's most valuable profession is that of a programmer archaeologist, because essentially, the layers of

software in all the large systems are just deep, interpenetrating, idiosyncratic, and interdependent. It's just become impossible to rewrite code for simplicity's sake. In fact, you can't replace the code that's already out there without wrecking the foundations of civilization.

There's this new job title or job path, the programmer archaeologist, who spends their time churning through this maddening nest of ancient languages looking for hidden and forgotten tools in our code to repair the existing programs or to find odd things that can be turned into unanticipated uses. In fact, in this new world, no one person writes code. It's teams. In fact, software development is a tradition that spans thousands of years, undertaken by these rotating shifts of programmers, these generational rotating shifts of programmers, 'cause they work for a while. Then they go in cold sleep. The next set come on, onwards and onwards, 'cause that's what it takes to understand and make changes in the future.

If you can imagine you just have layers and layers of scripts, and code, and hack, and APIs, and workarounds so deeply embedded into every piece of software that runs our everyday, that that's why this becomes a lifelong profession. Imagine that it was driven not by commercial pragmatism, 'cause it wasn't about the sale of software. It's that this is so necessary for the day to day survival of running these machines and these systems that have to span light-years and centuries. We're not gonna be there in 2027, right? That's not there. I think we're along the path to that, because if you look at the media today, it's full of a ton of scary stories saying that we're short of STEM-educated individuals—so science, technology, education, and math—to be our next workforce. The media uses a bunch of big numbers and a lot of math, which I question most of it, but to make us really scared that we don't have enough folks. We need to focus and get folks more educated in this space.

When we're thinking about the code and the software we're gonna see in 2027, whether some of that is written by machines or some of that is just the result of ourselves, it takes a snippet of the code and wraps it in another language, and wraps it in another, and wraps it in another, because that's what the programmer on call understood how to do. We're going for efficiency's sake to patch

things up quickly and go along. How are we gonna be able to figure that out in the future? If we're somewhere between today and this vision of, well, it was set in the 12th millennium, so a little bit further out in the future, right? If we're somewhere in that space today when we're looking at our threatcasting efforts for 2027, I would suggest that you think about the evolution of software, the evolution of computing, and what potentially new career paths might have developed as a result. Thank you.

.....

Appendix 3

63

RESEARCH SYNTHESIS WORKBOOKS

After listening to the six-curated inputs each group, assigned to one speaker, synthesized what they heard and plotted data points accordingly. With each data point they carefully examined implications of this data point, if the implication was positive or negative, and any thoughts around what might be done to encourage the positive data point or mitigate the negative. The first twelve pages of this appendix contain the role up of all the groups' data points for each speaker. This was necessary for the threatcasting inputs. The second half of this appendix shows all the raw data for each group individually.

The information found in the following pages is raw data and has not been spell checked or edited in any manner.

Slot	1			
Speaker	Dr. Genevieve Bell			
#	Data Point	Implication	Positive or Negative?	What should we do?
1	AI is its own cultural category	We need to interrogate it just as we approach other cultural categories	It allows us to expose the intrinsic biases behind AI	Develop a protocol for these conversations
2	AI is another manifestation of what means to be human	Makes AI normal/natural; AI has a certain set of values, belief systems based on who made/raised it	It creates an additional layer of complexity; risk of disenfranchising groups; magnify biases	Force transparency on how we treat AI in ways not accounted for in other technologies
3	AI can have a subconscious/id	AI operates with implicit agendas/ tacit motives	Negative. The motives behind their decisions may not be self-evident	Build mechanisms that allow us to challenge/interrogate AI decisions
4	Culture matters when interpreting observations	AI needs to understand our culture to understand us		There has to be a bi-directional flow of cultural understanding
5	Algorithms are not necessarily AI's	If a system cannot account for culture, it is not AI. AI is not just technology.	Positive:	Interrogate AIs as technologies and as cultures, two different sets of interrogation.
6	AI reflects the worldview/biases of its creators	misinterpreting recommendations based on bias	negative	incorporate more perspectives
7		male input could predispose AI toward conflict		
		S Who is creating AI? Is Silicon Valley really representative of all the users of these tools? Similarly, does Silicon Valley have an appropriately broad perspective to see all the problems that AI could be used to solve?	negative	Need to get innovation and input into innovation from outside traditional spheres.
8	AI is actually a broad constellation of technologies -- many of which are extremely dissimilar from each other	Innovation could lead to many blind alleys and changed directions	negative	segment (split into subfields) our conversation about AI, understand what we're actually talking about and how it maps to our goals
9	Skinner (and other early AI innovators) conceived of actions as deterministic -- can predict output based on input	Drives thought toward a deterministic model, which encourages us to think about inputs and map to outputs, and to understand likely outcome before we leap	positive	Returning to this drive could be useful in avoiding bad/unpredicted outcomes
10		The systems at play are incredibly complex -- often extremely hard to understand & predict. Thus, this approach could be a trap/could provide a false sense of security.	negative	constantly check our assumptions and recognize the uncertainty and instability of the input environments in which we operate.
		Our legal system & concept of liability relies on free will and understandable decisions made by humans. How does this work if machines make decisions w/real world consequences and we don't understand those decisions	neg/pos	
11	Inventors of AI were reacting to/against Freud & messiness of EU psychology	Drove inventors to look for clean, predictable, "machine" solutions -- drives greater focus and innovation on cause and effect	positive	

		But we are increasingly convinced that we live in a probabilistic world -- not a deterministic one. Humans are frightened of disorder, even though we create disorder without realizing it. Losing sight of these consequences risks creating (bad) unintended consequences.	negative	dreams and probabilistic actions are the result of incredibly complex systems -- we are already creating systems capable of emergent behavior. If we ignore this potential, we have no control over what we create.
12	"Artificial" carries connotations that are considered negative in today's society	AI is a technology rooted in the mindset of the 1950s, which is a very different mindset and values framework than we have today	negative	Need to think about how to bring AI into alignment with our modern value systems and the threats & opportunities that we see.
13		Could cause fear, give a science-fictiony connotation. Could make it harder to engage w/policy because people perceive it largely as sci-fi	negative	Need to educate those who engage to look beyond facile implications and think about what it actually means
14	Humans have long told stories and myths about animation & bringing things to life.	AI embodies humanity's fascination with pushing our limits and going beyond what we can achieve. Suggests that we might not make entirely rational decisions concerning AI	negative	likely to not think through consequences and could do something stupid.
15			positive	might push further than otherwise possible b/c we believe anything is possible.
16		Feeds into our natural god complex -- all the stories are cautionary tales.		
17	Describes AI as "just another manifestation of what it means to be human"	Does it really make sense to think of AI as "human"? Should it be constrained by our own vision of ourselves, and why do we persist in anthropomorphizing it and thus limiting it...	negative	Broaden our perspective in innovation to think about applications that don't fit a "human" model
		Continuing to push innovation in AI could actually change what it means to be human by redefining what is possible.	neg/pos	invest in proactive analysis of the social/cultural/ethical implications of the societal changes we are creating
18	AI can have culture just like any other community of organisms	"children are a map of their parents" -- what is the culture that we are teaching our mechanical children?	neg/pos	Need to think carefully about how our culture might be perceived/learned from/warped by a distinct intelligence seeking to interact with it
19	AI can have culture just like any other community of organisms	Key is the points of friction between our culture and AI emergent cultures -- where they differ we have an opportunity to learn, and there is opportunity for great mischief	neg/pos	How are we "designing" their culture; what opportunities do they have to develop culture outside of our control? How will that interact with our own culture (s)?

Slot	2			
Speaker	Sam Harris			
#	Data Point	Implication	Positive or Negative?	What should we do?
1	We are not near the summit of intelligence	Unknown unknowns	Both	Planning and building a sustainable eco system in case of a catastophic event
2	We do not not know how long to create Super Intel AI safely	AI could become uncontrollable by humans which would cause the machines to become uncontrollable	Negative	Establish regulatory boundries/guidance and develop cooperation between AI community and neuroscience
3	Humans are incapable of establishing appropriate responses to AI	Unable to treat AI respectfully	Negative	Establish education curriculum at an early age
4	There is an unregulated race to create the first Super Intel AI	Power exploitation	Negative	Establish treaties (checks and balances)
5	AI could be more intelligent than Humans which could create catastrophc consequences/circumstance	AI could make decisions that are not within the Human race best interest	Negative	Run simulations of potential scenarios to see outcomes and how to mitigate
6	Artificial Intelligence will destroy us	Society will dramatically change as a result of AI. Institutions will have to adapt. Current Governance policy models not responsive for technology change. We have no idea how to "do" AI safely	Negative. Could lead to an AI arms race	Begin developing policy framework
7	We are not near the pinnacle of intelligence	AI creates a percieved difference between our religious beliefs, morality, ethics, and personal beliefs and what we are creating with AI.	Both. Positive because it can bring disparate groups together. It can be negative because challenging widespread beliefs can induce conflict.	Begin exploring how AI interacts with various value systems (Islam, Communism, etc). Create potential guidelines for implementing AI within our own realm of influence.
8	We will continue to improve our smart machines	The consequences may outweigh the benefits of continued improvements of smart machines. At what point are we at diminishing returns	Both.	
9	He doesn't see how AI DOESN'T destroy us - Path to AI taking over has already happened	Process is irreversible	Sam Harris' perspective was that this is negative - we've already lost.	Build in oversight as much as possible now, but based on his argument how can we
10	We are NOT NEAR peak intelligence	AI systems may over take and over run humanity	Potentially negative without appropriate oversight, checks, and balances	His assumption that intelligence is equivalent to volume and speed of processing data or inputs leads to his assertions. This may not be an accurate understanding of intelligence
11	What happens when the AI systems develop themselves and we the humans do not notice	AI systems will take over without human awareness, and at that point it will be too late to apply effective controls	Again - his entire presentation had a cautionary and negative spin, without oversight, we become the ants to the exponentially more powerful AI systems	He does not provide any hopefull suggestions, however, we can attempt to focus on appropriate controls, checks, and balances - at least as much as we focus on functional improvements and capabilities
12	There is a possibility of a few obtaining full control of the value of AI systems, and further exacerbating societal and economic inequalities	AI systems become another vector for the privileged few to dominate, control, and exploit the worlds population	This is presented as a cautionary negative, the worst of humanity exploiting the worst aspects of technological advances	He offers no hope, however, appropriate controls could limit this effect, concepts of open source distribution of the technical knowledge could level the playing field

13	The only potentially positive outcome discussed in his presentation of an AI future is what do the humans do - one big burning man event, a life of luxury?	As has happened over the last 100 years with technological process, humans have gone from labor and agriculture to office jobs, if this keeps happening, what will humans do - degrade into a destructive behavior, or positive - how will we culturally adjust to more free time	This could go either way, largely driven by human social and psychological tendencies, a life like the Jetsons - or burning man	Begin now preparing for the societal and cultural adjustments that will come with greater automation of all aspects of human life
14	One possible way to control the inevitable "Terminator" future could be to fuse or imbed AI with humans in a Bio-Fused manner	Fusion of humans and AI at a fundamental level may provide some degree of control over what Sam Harris views as a runaway train	Potentially positive - some hope for ensuring a non-fatal outcome for humanity under world domination of AI	Rather than letting a few well-resourced for-profit companies chase this goal in the cowboy manner they are today - create a more controlled, regulated, and collaborative effort to ensure the outcome is positive, sustainable, and available to broader humanity
15	No idea how long it will take humanity to develop AI - SAFELY	Although it MAY be possible that the coming takeover of AI will be possible, Sam Harris does not know how long it will take humanity to figure out how to do it in a non-catastrophic manner	Potential hope - given humanity pursues AI with thoughtfulness and foresight	Focus on improving the knowledge, skills, and experience level of humans related to any aspect of systems, AI, automation, to ensure the appropriate quality checks, oversight, and regulated controls

Slot	3			
Speaker	DAVE GIOE			
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Cyber will augment, not replace, HUMINT	Increased Capability	Positive	Learn how to implement the technology most EFFECTIVELY
2	80/20 more important than ever	more detailed information more accessible	Negative	Extra layer of security in cover development, more intimate relationship with Case Officer for the developer
3	Convenience in cyber application can lead to inadvertent shortcuts	Next generation may not learn all the implicit basics due to technological aid	Negative	Develop programs that start at the basic, bare bones level and build up incrementally, slowly incorporating cyber.
4	Cleared individuals are at more risk than ever and it's getting worse.	FIS can specifically target the most sensitive and cleared individuals	Negative	Establish a new agency level organization with the sole mission of establishing, vetting, and protecting cleared individuals and their PII
5	HUMINT is eternally relevant	As long as humans, HUMINT will matter	Positive	Don't let emerging technologies overshadow the importance of HUMINT.
6	Cyber capabilities shape but do not replace HUMINT.	HUMINT will always have high value. Cyber aids the recruitment cycle and can be used to provide more potential high-quality targets for HUMINT, and also to identify their potential vulnerabilities	Positive	Maximize HUMINT effectiveness using cyber. Deploy AI to find more targets, more weaknesses, and then act as an assistant to the handler by watching assets, summarizing their activities, and flagging potential issues.
7	We all leave a digital data trail, and that trail is going to be filled with richer and richer data as we deploy more sensors and more computing into the world. IOT, wearables, analytics, etc.	Cyber analysis of everyone's data trail now makes it harder to maintain cover. As more sensors and more computing capability is deployed, maintaining convincing cover gets harder and harder.	Negative	Use digital means, including AI, to support a digital misinformation campaign (e.g. edit you into photos/videos of your "friends" in your cover story) and spoof a digital trail that is consistent with your cover story. Deploy AI to find holes in cover stories and flag them to handlers so that action can be taken to mitigate.
8	OSMINT and SOCMINT (social media intelligence) will matter more in the future than they do now. More information will be available on people as they spend more of their lives online in social media platforms. Social media will evolve rapidly to add virtual/augmented reality, voice and gesture. Connections will continue to expand beyond "friend-to-friend" to connect people to services, markets (a full-fledged transactional platform), businesses, and many other organizations.	OSMINT will become equal in value to HUMINT/IMGINT/SIGINT but not surpass it. Need to plan for value of OSINT/SOCMINT in the future, and invest now ready to harvest that in the future.	Positive	Build increased SOCMINT capability using big data and sophisticated AI-enhanced analytics tools. And assure continued access to social media platforms despite likely increase in encryption. Meld SOCMINT with HUMINT by ensuring HUMINT approaches are used inside social media environments.

9	Reliable encryption of communication now commonplace.	Safe, secure communication is much easier between handlers and their assets. But even with higher fidelity electronic communications (video-conferencing, and beyond that virtual reality/augmented reality conferencing experiences) nothing replaces the quality of communications that occurs during face-to-face contact. Good intelligence officers will pick up on the smallest non-verbal communications and use them to build hunches that prove to be vital.	Positive	Face-to-face communication is best for establishing trust. Once trust is established, leverage encrypted communications to boost the flow of information and maintain good frequency of communication with the asset. This should lead to a better flow of quality information. But never make electronic communication the only modality. Always maintain face-to-face communications as a way to look for other signals.
10	Major leaks (wikileaks, Snowden) hurt trust between handlers and agents.	Leaks undermine the trust that is vital for HUMINT. Every time a leak occurs, assets will fear that the next leak will expose them and put them at high risk, which may be existential.	Negative	Ephemeral communications (somewhat like snapchat) may prove to be an aid trust. If communication can be proven/guaranteed not to be recorded in any way, or is fully anonymised, would that help to boost trust? Also, brief all assets that leaks do/will occur, and that a quick-response mitigation is committed if a leak occurs that will extract the asset quickly and get them to a secure location.

Slot	4			
Speaker	Paul Thomas			
#	Data Point	Implication	Positive or Negative?	What should we do?
1	opportunity casting vs threatcasing	does this require a re-look at economic and risk models	Positive and negative ()	-increase high risk/high return research (gov resource allocation - or incentives for private)
2	risk may pay off	opportunity cost for not doing something	positive	encourage smaller lower cost decisions (flip it game theory)
3	appearance of new markets	new markets will form when there is a need	positive (as long as the market can be allowed or encouraged to form)	incentivize and encourage market development and shifts
4	power of coalitions	cooperation beyond disagreement in a balkanized world	positive overall (negative for weak/non-members)	rebuild / relook coalitions (willingness to adjust)
5	traditionally US being big giant in economic/trade relations (this is/will)	less global influence	negative	rebuild / relook coalitions
6	appropriately assessing risk	cyber security risk can be quantified	positive	encourage clear definitions and liabilities (understanding exposure, hazard and vuln)
7	Focusing on negative creates the possibility to ignore the positive	There are opportunities for alternate perspectives with positive implications	negative	Create "Opportunity casting" events
8		We are limiting our sample set due to focusing on the negative	negative	invert the sample collection to force awareness of other data points and sources
		If others are focusing on the positive implications they may work to accelerate movement toward something we see as a negative even though they are not malicious actors	negative (not bad, just that there are now groups working against us)	Look for the groups that benefit from your risk and what actions they are likely to take (put an opposing force mindset on and
9	Nations (or other groups) can take risks that individuals can't and vice versa	When a nation accepts a risk, it cascades to an individual. Vice Versa is can also be true	positive	we need to find ways to explore laws that are correct for a nation but create unacceptable risk as an individual
10		There is a risk of making the wrong choices as a nation because we are making the choices based on what is right for an individual	negative	we need to find ways to explore laws that are correct for a nation but create unacceptable risk as an individual
11	Once you start having gains against other players you open yourself up to interesting behavior like collations	This is not a zero-sum game	Negative	Model the coalitions that will form in opposition to your expected future - expected utility models can partly predict the future based on a few limited inputs



ASU
Interdis
Science
Technol
Building
Marston
Exploration
Gallery of
Scientific Ex

ISTB4
781
TERRACE

Slot	5			
Speaker	Andre LeBlanc			
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Eventually you will take an 8 year degree to learn what happened in 1 month	Standard career progression and societal value of 'degree' program becomes obsolete; Further specialization and hybrid education becomes the norm	Negative	Develop hybrid education system; Emphasize areas of specialty and skillsets in education
2	Technology will be a means to enhance the human experience, not hinder it	An average individual will understand the outputs of education and technology, not the inputs or mechanics behind it; it will be a largely cultural change	Yes	Ensure that there are people who specialize in one or the other - humanity cannot move forward without understanding an inherent knowledge of the mechanics behind technological occurrences (I.e. do not let a computer filter your life entirely for you)
3	AI will facilitate modeling/simulation of the medical field (specific cancers and drugs) through virtualization	AI will explode the medical field; Diseases will become far less serious because of virtualization; Diseases w/outliers identified compile results from experiments as a mass of knowledge; A system can be modeled to the 'n'th degree and totally analyzed and simulated by AI; Aspects of society can be virtualized - populations can be modeled and anticipated (businesses will control/manipulate underlying segments of society)	Positive	Avoidance of biometrics (if you hack a facial scan or fingerprint, it is lost forever because you can't change it); Begin protection programs and ubiquitous availability to all humans to counter it - the only way to counter something that is inevitable is to ensure that everyone has access to it (open source information)
4	By 2045 computers will have more intelligence than all humans combined	Humans will no longer be able to think holistically about what knowledge is; Identity crisis of what sets humans apart from other beings or machines on the planet	Yes	We can't do anything, it's inevitable; Every generation takes on a responsibility to pass on specific knowledge
5	Human mind can't keep up with data flow	Intelligence vs. knowledge; we need AI to work with and access the data flow; AI becomes the filter and enabler for our information	Yes	Ensure that robots will work for us; Morph the progression of technology/AI in order to act as a service, not as a sentient being
6	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.	The machine will not need to rely on a human operator to input an objective operator in to it. Intuitiveness is the gateway to Adaptivity and Self-learning. All three are interrelated. Current AI lacks intuition. Once this limit is removed it can identify what's important and what's not.	Currently a good thing that computers can only do the things for which we input instructions.	If your goal is a general purpose AI, (completely benevolent) you need to think of a new programming paradigm that does not rely on a human baking in the objective from the beginning. If a pessimist, you should limit AI's intuitiveness.
7	By 2035 AI will equal a SINGLE human's intelligence, by 2045 it will exceed ALL of human intelligence.	Could lead to a "Cyberdine Systems Skynet" scenario or an open world information library- "Google Docs on Steroids."	Both, depending upon whether or not it is weaponized and used to create false information or is used to expand and synchronize big data into better solutions to problems.	There needs to be transparency where the construction of a reliable AI mechanism can track AI activity to ensure the validity of AI decisions, simulations, etc. (checks and balances).

8	Moore's Law- (As of 2017) human knowledge doubles every 12 months and the human brain cannot comprehend all of this information.	Eventually, without systems that can bring information back to humans in a form they can understand, it becomes a Pandora's box. We will be forced to trust the judgement of some entity whose judgement we cannot fathom.	Explicitly neutral because it is completely divorced from what we as humans do with it, otherwise this data point is simply a philosophical exercise in epistemology.	Malevolence or benevolence needs to be manifested in action and not factual condition; exponential growth of information in and of itself is not enough. We can however, postulate what humans would do with the ability to exploit and/or weaponize the the exponential growth of human knowledge and its ability to be accessed at tremendous efficiency.
9	Machine self-adaption	No human in the loop that means self-adaption will lead to systems redefining parameters/objectives	Negative, without Command/Control there can be emergent AI behavior	"Big red button" (firewalls/hardstops). Question: Does AI learn forwards and backwards? Should we purposely use antiquated software/mechanical systems as purposeful segmenting of AI from critical systems?
10	AI replacing work	Jobs will be displaced; social unrest; social revolution?	(not positive/negative) No stagnation of society, relative wealth of population groups will not be static. Society "heat death"	Responsible innovation, examine 2nd/3rd order effects. (natural gas industry would provide recommendations on the decline/displacement of coal industry)
11	Virtual humans	Optimized drug treatments (which may not be effective if the virtual is not 100% accurate)	Negative, what % knowledge of understanding of the human body is good enough? Potentially dangerous situation where 99.95% may not be good enough. Positive, 'building' a human from the genome level in cyber world would allow infinite iterations of human development and further social development.	Ethical considerations, would a virtual human have agency? Does 'it' have rights? At the lowest building blocks of human development, how do we know when we have 100% knowledge of elements and mechanisms? As we develop a virtual human (and build an environment to simulate societal impacts on human development), does this universe have equal rights to existence as the ones humans exist in?

Slot	6			
Speaker	Natalie Vanatta			
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Progression from IoT to Internet of Everything. Moores law begins to break down due to laws of physics limitations. By 2027, microprocessors are estimated to be 5 nanometers thus the tendency towards embedded systems is poised to continue.	1. Big Data pertinence and pervasiveness will continue to increase. Secrecy and privacy concerns will increase. 2. The physical domain will become intertwined with the virtual domain. Military focus on kinetic operations will be challenged as cyber becomes more integrated with the physical world.	Both	- adjust TTPs to operate in this type of world (i.e. more information operations planning integrated into cyber)
2	Challenged government ability to conduct secret (covert, clandestine) operations in 2027?	(Includes friendly and enemy actions) Intelligence operations will be more closely married to information operations: signature based detection will go away, while enemies will be more deliberate in evading behavioral based intelligence capabilities.	Both	- adjust TTPs to operate in this type of world (i.e. more information operations planning integrated into cyber)
3	Converting vast amounts of data into actionable decision making and making decisions at the speed of light.	Machine learning, data science, artificial intelligence, and talent management becomes more important as data volumes grow exponentially.	Both	- Grow fields of data science to be more encompassing (bring in the computer scientists)
4	A need for programming archaeologists (and other new career fields).	Given complexity and vastness of code, human-in-the loop adaptations are best suited for career paths going forward. Human abilities must be augmented by machines. Focus on developing code to understand code. New careers like "data scientist" are already coming on board, WRT to humans parsing and understanding data with the help of machines and code. In the future, we can expect more jobs like this, but also expect the humanities side to catch up with the tech and call for expertise in dealing cyber ethics and other similar issues.	Both	- National objective to increase education and offset job loss due to automation
5	Collapse of Moore's Law	Physics will limit computing capabilities until we find an alternative.	Negative because it means we need alternative means to increase computing capacity.	Ensure we have the technological next step ready for when we hit the limits of physics.
6	Sensors are everywhere	With sensors everywhere, it will be more difficult (impossible?) to conduct covert and clandestine operations.	Positive if we are trying to see what is going on, negative if we are trying to hide.	Set specific guidelines on who controls the sensors, where they are and aren't allowed, etc.
7	Surrounded by computer power	With more computer power we will have more leisure - we need to know how to use that leisure time. Do we get to the point of "the internet of nothing" because it's been so devalued? Legacy ownership has changed.	Positive for the most part - it allows for greater computing overall, and an evolution of the Internet of Things.	We need to create a specific path forward into the future of ubiquitous computing. We can look at historical precedents for ideas.

8	AI in the data - processing - information - analysis - knowledge - judgment - situational understanding - decision process	Where is the division of labor? What parts of the cycle are better performed by humans and what parts are better performed by AI?	Positive (with caveats)	Determine how much power we want an AI to have in our decision-making process. Set a specific point where a human must be in the loop. Precedent already set with autonomous weapons - similar limitations should be set with all AI systems.
9	We become the machine / processor	Self-knowledge - will the people know that they are part of the processing power? / What is the social contract between groups? / Does this create an imagination deficiency? / Who controls the processing power that is being done by people? / Why are we doing all this? People will want to know why they're being used. / What is the Luddite response to humans as processors? Will there be a revolution or civil war? / There will be a change in power dynamics based on mass and computing power.	Negative if done incorrectly, (and there are way more ways for it to go wrong than right).	Create systems that ensure that humans are not abused by the system, whether by other humans or computers.
10	STEM deficiency	Do we need to train more arts?		
11	Programmer archaeologist?	Who are we serving?		
12	In 10 years, the width of transistors will be 5 atoms... Moore's Law will fail according to physics	Everything can be a sensor (IoT to loET)	Positive is that sensors will be everywhere; negative is that sensors will be everywhere	Assume everything will be a sensor
13	Today's iPhone has more computing power than was used to send a man to the moon	Continued exponential growth will impact the way we plan for the future	both	Assume pace will accelerate
14	Current Operations are conducted (1) Normal Operations; (2) Covert Operations; (Clandestine Operations	What data will be available for IPB?	both	Assume transparency will prevail; seek alternatives to covert/ clandestine operations
15	Will we continue to be able to conduct covert and/or clandestine operations in the future? (Given that sensors will be everywhere and transparency will become dominant)	Does it matter? Will we need to be much better at "normal" operations if covert and clandestine will become less impactful?	both	Assume transparency will prevail; seek alternatives to covert/ clandestine operations
16	Decision making must be performed at machine speed	What intelligence collection/analysis/processing/dissemination tasks are best performed by human interface and which ones are more amenable to AI?	N/A	Focus academia/industry/govt on R&D topics that will impact the TCEPD cycle. The amount of data available will be overwhelming.
17	How do we get decision makers the 10,000 hours of practice to achieve expertise in a given field?			
18	Will we need "Programmer Archaeologists" to help modify SW code to keep pace with AI?	Some code will be written by AI, but day-to-day survival may be dependant upon SW development teams		
19	Today we use humans to und	Data to information to knowledge to situational understanding requires an understanding of "why"		
20	Is the emphasis on STEM rel	In the future, emphasis on STEM education may be less important given advancement in AI		

Research Synthesis Worksheet				
Slot / Speaker	Dr. Genevieve Bell			
Group 1				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	AI is its own cultural category	We need to interrogate it just as we approach other cultural categories	It allows us to expose the intrinsic biases behind AI	Develop a protocol for these conversations
2	AI is another manifestation of what means to be human	Makes AI normal/natural; AI has a certain set of values, belief systems based on who made/raised it	It creates an additional layer of complexity; risk of disenfranchising groups; magnify biases	Force transparency on how we treat AI in ways not accounted for in other technologies
3	AI can have a subconscious/id	AI operates with implicit agendas/ tacit motives	Negative. The motives behind their decisions may not be self-evident	Build mechanisms that allow us to challenge/interrogate AI decisions
4	Culture matters when interpreting observations	AI needs to understand our culture to understand us		There has to be a bi-directional flow of cultural understanding
5	Algorithms are not necessarily AI's	If a system cannot account for culture, it is not AI. AI is not just technology.	Positive:	Interrogate AIs as technologies and as cultures, two different sets of interrogation.

Research Synthesis Worksheet				
Slot / Speaker	Sam Harris			
Group2				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	We are not near the summit of intelligence	Unknown unknowns	Both	Planning and building a sustainable eco system in case of a catastrophic event
2	We do not know how long to create Super Intel AI safely	AI could become uncontrollable by humans which would cause the machines to become uncontrollable	Negative	Establish regulatory boundaries/guidance and develop cooperation between AI community and neuroscience
3	Humans are incapable of establishing appropriate responses to AI	Unable to treat AI respectfully	Negative	Establish education curriculum at an early age
4	There is an unregulated race to create the first Super Intel AI	Power exploitation	Negative	Establish treaties (checks and balances)
5	AI could be more intelligent than Humans which could create catastrophic consequences/circumstance	AI could make decisions that are not within the Human race best interest	Negative	Run simulations of potential scenarios to see outcomes and how to mitigate

Research Synthesis Worksheet				
Slot / Speaker	DAVE GIOE			
Group 3				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Cyber will augment, not replace, HUMINT	Increased Capability	Positive	Learn how to implement the technology most EFFECTIVELY
2	80/20 more important than ever	more detailed information more accessible	Negative	Extra layer of security in cover development, more intimate relationship with Case Officer for the developer
3	Convenience in cyber applicaiton can lead to inadvertant shortcuts	Next generation may not learn all the implicit basics due to technological aid	Negative	Develop programs that start at the basic, bare bones level and build up incrementally, slowly incorporating cyber.
4	Cleared individuals are at more risk than ever and it's getting worse.	FIS can specifically target the most sensitive and cleared individuals	Negative	Establish a new agency level organization with the sole mission of establishing, vetting, and protecting cleared individuals and their PII
5	HUMINT is eternally relevant	As long humans, HUMINT will matter	Positive	Don't let emerging technologies overshadow the importance of HUMINT.

Research Synthesis Worksheet				
Slot / Speaker	Paul Thomas			
Group 4				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	opportunity casting vs threatcasing	does this require a re-look at economic and risk models	Positive and negative ()	-increase high risk/high return research (gov resource allocation - or incentives for private)
2	risk may pay off	there is an opportunity cost for not doing something	positive	encourage smaller lower cost decisions (flip it game theory - rivest)
3	appearance of new markets	new markets will form when there is a need	positive (as long as the market can be allowed or encouraged to form)	incentivize and encourage market development and shifts
4	power of coalitions	cooperation beyond disagreement in a balkanized world	positive overall (negative for week/non-members)	rebuild / relook coalitions (willingness to adjust)
5	traditionally US being big giant in economic/trade relations (this is/will)	less global influence	negative	rebuild / relook coalitions
6	appropriately assessing risk to	cyber securty risk can be	positive	encourage clear defintions and liabilities (understanding exposure, hazard and vuln)

Research Synthesis Worksheet				
Slot / Speaker	Andre LeBlanc			
Group 5				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Eventually you will take an 8 year degree to learn what happened in 1 month	Standard career progression and societal value of 'degree' program becomes obsolete; Further specialization and hybrid education becomes the norm	Negative	Develop hybrid education system; Emphasize areas of specialty and skillsets in education
2	Technology will be a means to enhance the human experience, not hinder it	An average individual will understand the outputs of education and technology, not the inputs or mechanics behind it; it will be a largely cultural change	Yes	Ensure that there are people who specialize in one or the other - humanity cannot move forward without understanding an inherent knowledge of the mechanics behind technological occurrences (I.e. do not let a computer filter your life entirely for you)
3	AI will facilitate modeling/simulation of the medical field (specific cancers and drugs) through virtualization	AI will explode the medical field; Diseases will become far less serious because of virtualization; Diseases w/outliers identified compile results from experiments as a mass of knowledge; A system can be modeled to the 'n'th degree and totally analyzed and simulated by AI; Aspects of society can be virtualized - populations can be modeled and anticipated (businesses will control/manipulate underlying segments of society)	Positive	Avoidance of biometrics (if you hack a facial scan or fingerprint, it is lost forever because you can't change it); Begin protection programs and ubiquitous availability to all humans to counter it - the only way to counter something that is inevitable is to ensure that everyone has access to it (open source information)
4	By 2045 computers will have more intelligence than all humans combined	Humans will no longer be able to think holistically about what knowledge is; Identity crisis of what sets humans apart from other beings or machines on the planet	Yes	We can't do anything, it's inevitable; Every generation takes on a responsibility to pass on specific knowledge
5	Human mind can't keep up with data flow	Intelligence vs. knowledge; we need AI to work with and access the data flow; AI becomes the filter and enabler for our information	Yes	Ensure that robots will work for us; Morph the progression of technology/AI in order to act as a service, not as a sentient being

Research Synthesis Worksheet				
Slot / Speaker	Vanatta			
Group 6				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Progression from IoT to Internet of Everything. Moores law begins to break down due to laws of physics limitations. By 2027, microprocessors are estimated to be 5 nanometers thus the tendency towards embedded systems is poised to continue.	1. Big Data pertinence and pervasiveness will continue to increase. Secrecy and privacy concerns will increase. 2. The physical domain will become intertwined with the virtual domain. Military focus on kinetic operations will be challenged as cyber becomes more integrated with the physical world.	Both	- adjust TTPs to operate in this type of world (I.e. more information operations planning integrated into cyber)
2	Challenged government ability to conduct secret (covert, clandestine) operations in 2027?	(Includes friendly and enemy actions) Intelligence operations will be more closely married to information operations: signature based detection will go away, while enemies will be more deliberate in evading behavioral based intelligence capabilities.	Both	- adjust TTPs to operate in this type of world (I.e. more information operations planning integrated into cyber)
3	Converting vast amounts of data into actionable decision making and making decisions at the speed of light.	Machine learning, data science, artificial intelligence, and talent management becomes more important as data volumes grow exponentially.	Both	- Grow fields of data science to be more encompassing (bring in the computer scientists)
4	A need for programming archaeologists (and other new career fields).	Given complexity and vastness of code, human-in-the loop adaptations are best suited for career paths going forward. Human abilities must be augmented by machines. Focus on developing code to understand code. New careers like "data scientist" are already coming on board, WRT to humans parsing and understanding data with the help of machines and code. In the future, we can expect more jobs like this, but also expect the humanities side to catch up with the tech and call for expertise in dealing cyber ethics and other similar issues.	Both	- National objective to increase education and offset job loss due to automation

Research Synthesis Worksheet				
Slot / Speaker	bell			
Group 7				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	AI reflects the worldview/biases of its creators	misinterpreting recommendations based on bias	negative	incorporate more perspectives
2		male input could predispose AI toward conflict		
		Who is creating AI? Is Silicon Valley really representative of all the users of these tools? Similarly, does Silicon Valley have an appropriately broad perspective to see all the problems that AI could be used to solve?	negative	Need to get innovation and input into innovation from outside traditional spheres.
3	AI is actually a broad constellation of technologies -- many of which are extremely dissimilar from each other	Innovation could lead to many blind alleys and changed directions	negative	segment (split into subfields) our conversation about AI, understand what we're actually talking about and how it maps to our goals
4	Skinner (and other early AI innovators) conceived of actions as deterministic -- can predict output based on input	Drives thought toward a deterministic model, which encourages us to think about inputs and map to outputs, and to understand likely outcome before we leap	positive	Returning to this drive could be useful in avoiding bad/unpredicted outcomes
5		The systems at play are incredibly complex -- often extremely hard to understand & predict. Thus, this approach could be a trap/could provide a false sense of security.	negative	constantly check our assumptions and recognize the uncertainty and instability of the input environments in which we operate.
		Our legal system & concept of liability relies on free will and understandable decisions made by humans. How does this work if machines make decisions w/real world consequences and we don't understand those decisions	neg/pos	
5.5	Inventors of AI were reacting to/against Freud & messiness of EU psychology	Drove inventors to look for clean, predictable, "machine" solutions -- drives greater focus and innovation on cause and effect	positive	
		But we are increasingly convinced that we live in a probabilistic world -- not a deterministic one. Humans are frightened of disorder, even though we create disorder without realizing it. Losing sight of these consequences risks creating (bad) unintended consequences.	negative	dreams and probabilistic actions are the result of incredibly complex systems -- we are already creating systems capable of emergent behavior. If we ignore this potential, we have no control over what we create.
6	"Artificial" carries connotations that are considered negative in today's society	AI is a technology rooted in the mindset of the 1950s, which is a very different mindset and values framework than we have today	negative	Need to think about how to bring AI into alignment with our modern value systems and the threats & opportunities that we see.
7		Could cause fear, give a science-fiction connotation. Could make it harder to engage w/policy because people perceive it largely as scifi	negative	Need to educate those who engage to look beyond facile implications and think about what it actually means

8	Humans have long told stories and myths about animation & bringing things to life.	AI embodies humanity's fascination with pushing our limits and going beyond what we can achieve. Suggests that we might not make entirely rational decisions concerning AI	negative	likely to not think through consequences and could do something stupid.
9			positive	might push further than otherwise possible b/c we believe anything is possible.
10		Feeds into our natural god complex -- all the stories are cautionary tales.		
11	Describes AI as "just another manifestation of what it means to be human"	Does it really make sense to think of AI as "human" Should it be constrained by our own vision of ourselves, and why do we persist in anthropomorphizing it and thus limiting it...	negative	Broaden our perspective in innovation to think about applications that don't fit a "human" model
		Continuing to push innovation in AI could actually change what it means to be human by redefining what is possible.	neg/pos	invest in proactive analysis of the social/cultural/ethical implications of the societal changes we are creating
12	AI can have culture just like any other community of organisms	"children are a map of their parents" -- what is the culture that we are teaching our mechanical children?	neg/pos	Need to think carefully about how our culture might be perceived/learned from/warped by a distinct intelligence seeking to interact with it
13		Key is the points of friction between our culture and AI emergent cultures -- where they differ we have an opportunity to learn, and there is opportunity for great mischief	neg/pos	How are we "designing" their culture; what opportunities do they have to develop culture outside of our control? How will that interact with our own culture(s)?

Research Synthesis Worksheet				
Slot / Speaker	Sam Harris			
Group 8				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Artificial Intelligence will destroy us	Society will dramatically change as a result of AI. Institutions will have to adapt. Current Governance policy models not responsive for technology change. We have no idea how to "do" AI safely	Negative. Could lead to an AI arms race	Begin developing policy framework
2	We are not near the pinnacle of intelligence	AI creates a perceived difference between our religious beliefs, morality, ethics, and personal beliefs and what we are creating with AI.	Both. Positive because it can bring disparate groups together. It can be negative because challenging widespread beliefs can induce conflict.	Begin exploring how AI interacts with various value systems (Islam, Communism, etc). Create potential guidelines for implementing AI within our own realm of influence.
3	We will continue to improve our smart machines	The consequences may outweigh the benefits of continued improvements of smart machines. At what point are we at diminishing returns	Both.	

Research Synthesis Worksheet				
Slot / Speaker	Dr David Gioe			
Group 9				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Cyber capabilities shape but do not replace HUMINT.	HUMINT will always have high value. Cyber aids the recruitment cycle and can be used to provide more potential high-quality targets for HUMINT, and also to identify their potential vulnerabilities	Positive	Maximize HUMINT effectiveness using cyber. Deploy AI to find more targets, more weaknesses, and then act as an assistant to the handler by watching assets, summarizing their activities, and flagging potential issues.
2	We all leave a digital data trail, and that trail is going to be filled with richer and richer data as we deploy more sensors and more computing into the world. IOT, wearables, analytics, etc.	Cyber analysis of everyone's data trail now makes it harder to maintain cover. As more sensors and more computing capability is deployed, maintaining convincing cover gets harder and harder.	Negative	Use digital means, including AI, to support a digital misinformation campaign (e. g. edit you into photos/videos of your "friends" in your cover story) and spoof a digital trail that is consistent with your cover story. Deploy AI to find holes in cover stories and flag them to handlers so that action can be taken to mitigate.
3	OSMINT and SOCMINT (social media intelligence) will matter more in the future than they do now. More information will be available on people as they spend more of their lives online in social media platforms. Social media will evolve rapidly to add virtual/augmented reality, voice and gesture. Connections will continue to expand beyond "friend-to-friend" to connect people to services, markets (a full-fledged transactional platform), businesses, and many other organizations.	OSMINT will become equal in value to HUMINT/IMGINT/SIGINT but not surpass it. Need to plan for value of OSINT/SOCMINT in the future, and invest now ready to harvest that in the future.	Positive	Build increased SOCMINT capability using big data and sophisticated AI-enhanced analytics tools. And assure continued access to social media platforms despite likely increase in encryption. Meld SOCMINT with HUMINT by ensuring HUMINT approaches are used inside social media environments.
4	Reliable encryption of communication now commonplace.	Safe, secure communication is much easier between handlers and their assets. But even with higher fidelity electronic communications (video-conferencing, and beyond that virtual reality/augmented reality conferencing experiences) nothing replaces the quality of communications that occurs during face-to-face contact. Good intelligence officers will pick up on the smallest non-verbal communications and use them to build hunches that prove to be vital.	Positive	Face-to-face communication is best for establishing trust. Once trust is established, leverage encrypted communications to boost the flow of information and maintain good frequency of communication with the asset. This should lead to a better flow of quality information. But never make electronic communication the only modality. Always maintain face-to-face communications as a way to look for other signals.
5	Major leaks (wikileaks, Snowden) hurt trust between handlers and agents.	Leaks undermine the trust that is vital for HUMINT. Every time a leak occurs, assets will fear that the next leak will expose them and put them at high risk, which may be existential.	Negative	Ephemeral communications (somewhat like snapchat) may prove to be an aid trust. If communication can be proven/guaranteed not to be recorded in any way, or is fully anonymised, would that help to boost trust? Also, brief all assets that leaks do/will occur, and that a quick-response mitigation is committed if a leak occurs that will extract the asset quickly and get them to a secure location.

Research Synthesis Worksheet				
Slot / Speaker	Paul Thomas			
Group 10				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Focusing on negative creates the possibility to ignore the positive	There are opportunities for alternate perspectives with positive implications	negative	Create "Opportunity casting" events
2		We are limiting our sample set due to focusing on the negative	negative	invert the sample collection to force awareness of other data points and sources
Look for the groups that benefit from your risk and what actions they are likely to take (put an opposing force mindset on and play through the theater perspective)		If others are focusing on the positive implications they may work to accelerate movement toward something we see as a negative even though they are not malicious actors	negative (not bad, just that there are now groups working against us)	Look for the groups that benefit from your risk and what actions they are likely to take (put an opposing force mindset on and trace the threat from the opposition's point of view as being a positive change for them)
5	Nations (or other groups) can take risks that individuals can't and vice versa	When a nation accepts a risk, it cascades to an individual. Vice Versa is can also be true	positive	we need to find ways to explore laws that are correct for a nation but create unacceptable risk as an individual
6		There is a risk of making the wrong choices as a nation because we are making the choices based on what is right for an individual	negative	we need to find ways to explore laws that are correct for a nation but create unacceptable risk as an individual
7	Once you start having gains against other players you open yourself up to interesting behavior like collations	This is not a zero-sum game	Negative	Model the coalitions that will form in opposition to your expected future - expected utility models can partly predict the future based on a few limited inputs
8		This is not a zero-sum game	Positive	we must explore alternate relationships with groups we have existing relationships with we have to be aware of the potential for groups who would normally never cooperate to do so to gain competitive advantage; we should also find ways to model possible coalition formations
9	You sometimes need to expose yourself to risk to gain credibility (WWI web of treaties - necessary to gain trust but with big consequences)	there may be groups that are pretending to commit to risk but are actually hedging their bets	Negative	we need to find a way to verify commitment level; be aware of how we ally ourselves and how our risks tie us together (treaties)
10	New markets occur all the time	Risks may be mitigated by an emerging market	Positive	Invest in risky markets if they have the potential to engineer our desired future
11		New markets may emerge that make things worse (HFT-High Frequency Trading, organ transplant black markets)	Negative	Plan/find methods for mitigating the consequences of the new markets causing greater (negative) impacts

Research Synthesis Worksheet				
Slot / Speaker	Andrew LeBlanc			
Group 11				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.	The machine will not need to rely on a human operator to input an objective operator in to it. Intuitiveness is the gateway to Adaptivity and Self-learning. All three are interrelated. Current AI lacks intuition. Once this limit is removed it can identify what's important and what's not.	Currently a good thing that computers can only do the things for which we input instructions.	If your goal is a general purpose AI, (completely benevolent) you need to think of a new programming paradigm that does not rely on a human being in the objective from the beginning. If a pessimist, you should limit AI's intuitiveness.
2	By 2035 AI will equal a SINGLE human's intelligence, by 2045 it will exceed ALL of human intelligence.	Could lead to a "Cyberdine Systems Skynet" scenario or an open world information library- "Google Docs on Steroids."	Both, depending upon whether or not it is weaponized and used to create false information or is used to expand and synchronize big data into better solutions to problems.	There needs to be transparency where the construction of a reliable AI mechanism can track AI activity to ensure the validity of AI decisions, simulations, etc. (checks and balances).
3	Moore's Law- (As of 2017) human knowledge doubles every 12 months and the human brain cannot comprehend all of this information.	Eventually, without systems that can bring information back to humans in a form they can understand, it becomes a Pandora's box. We will be forced to trust the judgement of some entity whose judgement we cannot fathom.	Explicitly neutral because it is completely divorced from what we as humans do with it, otherwise this data point is simply a philosophical exercise in epistemology.	Malevolence or benevolence needs to be manifested in action and not just a condition; exponential growth of information in and of itself is not enough. We can however, postulate what humans would do with the ability to exploit and/or weaponize the exponential growth of human knowledge and its ability to be accessed at tremendous efficiency.

Research Synthesis Worksheet				
Slot / Speaker	Natalie Vannata			
Group 12				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Collapse of Moore's Law	Physics will limit computing capabilities until we find an alternative.	Negative because it means we need alternative means to increase computing capacity.	Ensure we have the technological next step ready for when we hit the limits of physics.
2	Sensors are everywhere	With sensors everywhere, it will be more difficult (impossible?) to conduct covert and clandestine operations.	Positive if we are trying to see what is going on, negative if we are trying to hide.	Set specific guidelines on who controls the sensors, where they are and aren't allowed, etc.
3	Surrounded by computer power	With more computer power we will have more leisure - we need to know how to use that leisure time. Do we get to the point of "the internet of nothing" because it's been so devalued? Legacy ownership has changed.	Positive for the most part - it allows for greater computing overall, and an evolution of the Internet of Things.	We need to create a specific path forward into the future of ubiquitous computing. We can look at historical precedents for ideas.
4	AI in the data - processing - information - analysis - knowledge - judgment - situational understanding - decision process	Where is the division of labor? What parts of the cycle are better performed by humans and what parts are better performed by AI?	Positive (with caveats)	Determine how much power we want an AI to have in our decision-making process. Set a specific point where a human must be in the loop. Precedent already set with autonomous weapons - similar limitations should be set with all AI systems.
5	We become the machine / processor	Self-knowledge - will the people know that they are part of the processing power? / What is the social contract between groups? / Does this create an imagination deficiency? / Who controls the processing power that is being done by people? / Why are we doing all this? People will want to know why they're being used. / What is the Luddite response to humans as processors? Will there be a revolution or civil war? / There will be a change in power dynamics based on mass and computing power.	Negative if done incorrectly, (and there are way more ways for it to go wrong than right).	Create systems that ensure that humans are not abused by the system, whether by other humans or computers.
6	STEM deficiency	Do we need to train more arts?		
7	Programmer archaeologist?	Who are we serving?		

Research Synthesis Worksheet				
Slot / Speaker	Lawrence Vacek			
Group 13				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	He doesn't see how AI DOESN'T destroy us - Path to AI taking over has already happened	Process is irreversible	Sam Harris' perspective was that this is negative - we've already lost.	Build in oversight as much as possible now, but based on his argument how can we
2	We are NOT NEAR peak intelligence	AI systems may over take and over run humanity	Potentially negative without appropriate oversight, checks, and balances	His assumption that intelligence is equivalent to volume and speed of processing data or inputs leads to his assertions. This may not be an accurate understanding of intelligence
3	What happens when the AI systems develop themselves and we the humans do not notice	AI systems will take over without human awareness, and at that point it will be too late to apply effective controls	Again - his entire presentation had a cautionary and negative spin, without oversight, we become the ants to the exponentially more powerful AI systems	He does not provide any hopeful suggestions, however, we can attempt to focus on appropriate controls, checks, and balances - at least as much as we focus on functional improvements and capabilities
4	There is a possibility of a few obtaining full control of the value of AI systems, and further exacerbating societal and economic inequalities	AI systems become another vector for the privileged few to dominate, control, and exploit the world's population	This is presented as a cautionary negative, the worst of humanity exploiting the worst aspects of technological advances	He offers no hope, however, appropriate controls could limit this effect, concepts of open source distribution of the technical knowledge could level the playing field
5	The only potentially positive outcome discussed in his presentation of an AI future is what do the humans do - one big burning man event, a life of luxury?	As has happened over the last 100 years with technological process, humans have gone from labor and agriculture to office jobs, if this keeps happening, what will humans do - degrade into a destructive behavior, or positive - how will we culturally adjust to more free time	This could go either way, largely driven by human social and psychological tendencies, a life like the Jetsons - or burning man	Begin now preparing for the societal and cultural adjustments that will come with greater automation of all aspects of human life
Backup	One possible way to control the inevitable "Terminator" future could be to fuse or imbed AI with humans in a Bio-Fused manner	Fusion of humans and AI at a fundamental level may provide some degree of control over what Sam Harris views as a runaway train	Potentially positive - some hope for ensuring a non-fatal outcome for humanity under world domination of AI	Rather than letting a few well-resourced for-profit companies chase this goal in the cowboy manner they are today - create a more controlled, regulated, and collaborative effort to ensure the outcome is positive, sustainable, and available to broader humanity
	No idea how long it will take humanity to develop AI - SAFELY	Although it MAY be possible that the coming takeover of AI will be possible, Sam Harris does not know how long it will take humanity to figure out how to do it in a non-catastrophic manner	Potential hope - given humanity pursues AI with thoughtfulness and foresight	Focus on improving the knowledge, skills, and experience level of humans related to any aspect of systems, AI, automation, to ensure the appropriate quality checks, oversight, and regulated controls

Research Synthesis Worksheet				
Slot / Speaker	Andre LeBlanc			
Group 14				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	Machine self-adaption	No human in the loop that means self-adaption will lead to systems redefining parameters/objectives	Negative, without Command/Control there can be emergent AI behavior	"Big red button" (firewalls/hardstops). Question: Does AI learn forwards and backwards? Should we purposely use antiquated software/mechanical systems as purposeful segmenting of AI from critical systems?
2	AI replacing work	Jobs will be displaced; social unrest; social revolution?	(not positive/negative) No stagnation of society, relative wealth of population groups will not be static. Society "heat death"	Responsible innovation, examine 2nd/3rd order effects. (natural gas industry would provide recommendations on the decline/displacement of coal industry)
3	Virtual humans	Optimized drug treatments (which may not be effective if the virtual is not 100% accurate)	Negative, what % knowledge of understanding of the human body is good enough? Potentially dangerous situation where 99.95% may not be good enough. Positive, 'building' a human from the genome level in cyber world would allow infinite iterations of human development and further social development.	Ethical considerations, would a virtual human have agency? Does 'it' have rights? At the lowest building blocks of human development, how do we know when we have 100% knowledge of elements and mechanisms? As we develop a virtual human (and build an environment to simulate societal impacts on human development), does this universe have equal rights to existence as the ones humans exist in?

Research Synthesis Worksheet				
Slot / Speaker	MAJ Natalie Vanetta			
Group 15				
#	Data Point	Implication	Positive or Negative?	What should we do?
1	In 10 years, the width of transistors will be 5 atoms... Moore's Law will fail according to physics	Everything can be a sensor (IoT to IoET)	Positive is that sensors will be everywhere; negative is that sensors will be everywhere	Assume everything will be a sensor
2	Today's iPhone has more computing power than was used to send a man to the moon	Continued exponential growth will impact the way we plan for the future	both	Assume pace will accelerate
3	Current Operations are conducted (1) Normal Operations; (2) Covert Operations; (Clandestine Operations	What data will be available for IPB?	both	Assume transparency will prevail; seek alternatives to covert/clandestine operations
4	Will we continue to be able to conduct covert and/or clandestine operations in the future? (Given that sensors will be everywhere and transparency will become dominant)	Does it matter? Will we need to be much better at "normal" operations if covert and clandestine will become less impactful?	both	Assume transparency will prevail; seek alternatives to covert/clandestine operations
5	Decision making must be performed at machine speed	What intelligence collection/analysis /processing/dissemination tasks are best performed by human interface and which ones are more amenable to AI?	N/A	Focus academia/industry/govt on R&D topics that will impact the TCEPD cycle. The amount of data available will be overwhelming.
	How do we get decision makers the 10,000 hours of practice to achieve expertise in a given field?			
	Will we need "Programmer Archaeologists" to help modify SW code to keep pace with AI?	Some code will be written by AI, but day-to-day survival may be dependant upon SW development teams		
	Today we use humans to understand "why"...	Data to information to knowledge to situational understanding requires an understanding of "why"		
	Is the emphasis on STEM relevant?	In the future, emphasis on STEM education may be less important given advancement in AI		

Appendix 4

FUTURES WORKBOOK DAY ONE AND TWO

In groups, participants develop scenarios based on data inputs from each speaker. The inputs were randomly selected. These scenarios followed a strict outline designed to envision a person in a place with a problem. Participants answered a variety of questions about their character including, “Describe how your person experiences the threat.” In addition to designing future scenarios from an individual character’s perspective, groups also explored the experience of the adversary.

Finally, groups were pushed to backcast. This foresight tool defined – what we have control over, what we do not have control over, and steps we should take to disrupt, mitigate, and recover from these futures four and eight years out.

This exercise was done twice, once each day, and the workbooks were used to inform the scenarios, found in this report. Participants had between one to two hours to complete the threatcasting process.

The information found in the following pages is raw data and has not been spell checked or edited in any manner.

Group 1		
Experience Title:	"Workforce Optimization"	
Estimated Date:	2027	
Data Points		
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.		5 MINUTES
Slot #1	Humans have long told stories and myths about animation & bringing things to life.	
Slot #2	We are not near the summit of intelligence - Training the labor force w/ necessary skills is difficult; AI available can't just interact w/ anybody; it requires a level of skill that excludes most workers	
Slot #3	Major leaks (wikileaks, Snowden) hurt trust between handlers and agents. - Health and Engineering systems are sharing information	
Slot #4	appropriately assessing risk to enable insurance market	
Slot #5	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.	
Slot #6	Collapse of Moore's Law	
PART ONE: Who is your Person? Meriam		
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.		15 MINUTES
Who is your person and what is their broader community?	Meriam, Syrian immigrant since 2015; lost husband to an industrial accident in 2017; works at Hephaestus Industries (metal fabrication company); has chronic asthma	
Where do they live?	Lafayette, Louisiana	
What is the threat?	That the AI-driven workforce optimization/re-engineering process is inferring from its partner Health Insurance AI. Both AI's are trying to acquire information to make better decisions;	
Briefly describe how your person experiences the threat.		
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?		
What is the experience we want the person to have with the threat?		
What is the experience we want them to avoid?		
	She gets reallocated to a higher risk environment at the factory; she looks around and notices that all the medical "misfits" are working in the high risk areas- creates a tension with her medical case worker. We want transparency; we want her to know the reason for her assignment and avoid hidden biases; work conditions should not exacerbate her . Avoid placing them in an environment that will exacerbate her health conditions	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)		10 MINUTES
"The Event" - How will your person first hear about or experience the threat?		

What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	What is different and/or the same as previous events or instantiations of the threat?	
	Moved from clean manufacturing to shipping; notices that everybody around her has similar health conditions; people are sick a lot; AI is not trying to kill her- we want to get her sick so she gets fired/drops out	
Question Two	How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
	As she logs in using biometrics, she gets a notice that she's getting moved	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
	Human barrier; systems shouldn't be sharing information;	
Question Two	Business Models: What new business models and practices will be in place to enable the threat?	
	Unintended linkage between Workforce Optimization and Healthcare System; Non-parameterized workforce optimization	
PART FOUR- Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Isolation of distinct AI systems; non-interference of processes	
2	AI system that can be interrogated about its decisions	
3	Legal protections	
Flags:		

What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Reuse of health insurance records	Government (US Congress)
2	Social status- They don't have the same mobility	
3	Self-improving AI	Developer/Industry
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Change privacy laws to include protection against AI access	
2	Options for asking AI for the parameters it is using	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Enforcement of non-interference between AIs	

Group 213	
Experience Title:	Misplaced Trust
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	1 - AI is its own cultural category
Slot #2	3 - Humans are incapable of establishing appropriate responses to AI
Slot #3	4 - Cleared individuals are at more risk than ever and it's getting worse.
Slot #4	1 - opportunity casting vs threatcasting
Slot #5	7 - By 2035 AI will equal a SINGLE human's intelligence, by 2045 it will exceed ALL of human intelligence.
Slot #6	13 - Today's iPhone has more computing power than was used to send a man to the moon
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Lauren
Where do they live?	Charlotte, North Carolina
What is the threat?	Safety to child abduction
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve?	
What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
Lauren is a middle schooler that bikes everywhere within the local area. Her parents installed a virtual personal assistant (VPA) to her phone that manages every aspect of her life and informs her parents on her whereabouts. The VPA tracks her homework, diet, extracurricular activities, etc. A nefarious actor looking to exploit this vulnerability for capital gain hacks the app and posts her profile for sale on the dark web. A cartel buys her profile. Through this, they can track the behavior and location of Lauren. With this information, the cartel can manipulate the behavior of Lauren in minor adjustments over the course of years because they have access to the commands of the VPA and the advice that it gives to its users. Additionally, they can recognize when Lauren is most vulnerable so that they can divert her at the opportune moment to abduct her.	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	

When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	What is different and/or the same as previous events or instantiations of the threat?	
	The degree of integration with other data sources, sensor, all aspects of individuals lives	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	
	As the broader population becomes more and more dependent and trusting on such services or capabilities, exposure and risk increases significantly. This is a method of exploitation that can be used by terrorist organizations targeting disgruntled security employees, loan sharks to target financially unstable families, gangs to target young boys, cartels to target girls for human trafficking, and any other genre of human activity.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Business Models: What new business models and practices will be in place to enable the threat?	
	More full and complete integration and trust of AI supplemented services into daily life	
Question Two	Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
	Basic conveniences of technology - GPS, reminder apps, social apps, Internet of Things, Phone assistants like Siri, Ok Google	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Laws, regulations, oversight, accountability for security and penalties for non-compliance	Legal, Fed/State
2	AI based sensors, checks & balances to identify and detect unapproved access or activity	Industry
3	Integration of seamless, transparent, rigorous security controls for user access (bio/behavior)	Industry

Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Use of non-compliant (global/external) developed apps (other countries, different laws)	Global Coalition
2	Difficult to address criminal and malicious activity - it always adopts new/emerging tech.	Global Legal
3	Unable to control the attribution problem	Global Legal
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Global Coalition established to address digital AI codes of conduct, personal sovereignty, digital jurisdiction areas	
2	Infrastructure grid cooperation between different sensors and data farms	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Stranger danger for the digital world. Education on how to appropriately use a VPA (WALL-E can be Smoky the Bear)	
2	Global Cybersecurity taskforce/ digital interpol	
3	Commercial industry regulated by government policies that review the security levels of apps	

Group 93	
Experience Title:	Hacked Doctor
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	#14 Humans have long told stories and myths about animation & bringing things to life.
Slot #2	#4: There is an unregulated race to create the first Super Intel AI
Slot #3	#4: Cleared individuals are at more risk than ever and it's getting worse.
Slot #4	#11: Once you start having gains against other players you open yourself up to interesting behavior like coalitions
Slot #5	#6: By 2035 AI will equal a SINGLE human's intelligence, by 2045 it will exceed ALL of human intelligence.
Slot #6	#2: Challenged government ability to conduct secret (covert, clandestine) operations in 2027?
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Lei lives in Taiwan. She's 37. No kids. Married. Doctor.
Where do they live?	Taipei.
What is the threat?	Can't keep up with medical progress, and now relies on an AI to help her keep up with the latest medical techniques. Her AI is now at risk of being hacked, or be at risk of ransomware.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	AI becomes a trusted partner, and is then hacked...could be used to influence Lei and get her to do things she wouldn't normally do. She might not be willing to be without it. How do you signal to Lei when the AI might be lying to her (because it's been hacked).
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	

What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	
Question One	Paste Question HERE
	"The Event" - How will your person first hear about or experience the threat?
	She starts to suspect that her AI is lying to her because it suggests that she does something that doesn't quite fit with her training.. for example to give a drug to a patient that may have an allergy to it.
Question Two	Paste Question HERE
	What are the broader implications of a threat like this? What might a ripple effect look like?
	Lawyers could hack AI of doctors to create malpractice lawsuits.
	Motives: for money, to ruin the doctor's reputation, or to change her behavior to impact a third party.
	Influence groups of people in a coordinated effort to make them unwittingly perform a task for you
	Influence everyone's AI and "hack the truth"
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	
Questions (pick two)	10 MINUTES
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
Business Models: What new business models and practices will be in place to enable the threat?	
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?	
Question One	Paste Question HERE
	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?
	Cultural: make it more normal for people to be supported by an AI to make decisions
	Technological: need to understand the AI black box in order to influence it. Have to spoof the learning set over time to shift the machine learning algorithms.
	Spoofing AI to say believable things rather than things that are easy to spot as being wildly out of bounds (e.g the AI equivalent of the Nigerian prince...)
Question Two	Paste Question HERE
	Business Models: What new business models and practices will be in place to enable the threat?
	Dark web service offering access to an individual's AI
	Analysts that understand the constellation of people in a target's life, which ones have AIs that can be influenced, and then build a plan to use those assets to achieve a goal.
	Build plans to maintain trust but subvert truth slowly over time so as to avoid suspicion by remaining clandestine.
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)	
Examine the combination of both the Experience Questions as well as the Enabling Questions.	
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.	

Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
	1 AI that is assigned to watch the output of other AIs to watch for shifts in their behavior that would indicate an attack may have occurred.	
	2 AI backups to recover to last known clean state	
	3 Firewalls and multi-factor authentication for AIs	
	4 Blockchain as a way to establish roots of trust?	
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
	1 Lie detector is created for AIs, enabling HUMINT techniques to be applied to catch AIs in a lie.	
	2 AIs start to develop a conscience and reject bad data because they start to "know better"	
	3 New programming tools for AIs are developed that help developers to make AIs more secure	
	4 Malpractice insurance becomes harder to get as a doctor unless you have a certified AI	
	5 AI becomes a "must have" for a set of professions, or those practitioners aren't considered to be up to standard	
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
	1 New programming languages to build AIs	
	2 Root of trust (blockchain?) and high-quality authentication applied to AI updates	
	3 Research the problem landscape to better understand it	
	4 Training for people using AIs to help them watch for inconsistencies	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
	1 New guardian AI that watches over your AIs	
	2 Societal conversation about where the risk of AI is greater than the potential benefits	

Group 4		
Experience Title:	AI Garbage Collection	
Estimated Date:	2027	
Data Points		
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES	
Slot #1	(19) AI can have culture just like any other community of organisms	
Slot #2	(2) We do not not know how long to create Super Intel AI safely	
Slot #3	(7) We all leave a digital data trail, and that trail is going to be filled with richer and richer data as we deploy more sensors and more computing into the world. IOT, wearables, analytics, etc.	
Slot #4	(14) You sometimes need to expose yourself to risk to gain credibility (WWI web of treaties - necessary to gain trust but with big consequences)	
Slot #5	(2) Technology will be a means to enhance the human experience, not hinder it	
Slot #6	(13) Today's iPhone has more computing power than was used to send a man to the moon.	
PART ONE: Who is your Person?		
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES	
Who is your person and what is their broader community?	James St. James, white male, 55 years old, sanitation worker	
Where do they live?	Houston, Texas	
What is the threat?	Loss of livelihood as automated, self-driving sanitation collection systems have vastly reduced need for human employees.	
Briefly describe how your person experiences the threat.		
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve?		
What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?		
What is the experience we want the person to have with the threat?		
What is the experience we want them to avoid?		

	After introduction of automated, self-driving sanitation trucks in 2024, James now works half his former hours and has lost a long-anticipated promotion to supervisor, currently riding the automated truck as an emergency failsafe. He monitors the system as it follows its route, noting data such as inaccurate GPS locations or dumpsters at the wrong angle, anything for which the automated trash truck would not be able to account. His positions will be thoroughly redundant in one to three years as testing is completed and the system is certified fully operational. At that time, he will be laid off. He wants to avoid unemployment and would like to gain a full time management position but is competing against dozens of his former peers who have also been downsized. Working with the systems each day, and talking to his colleagues, he begins to realize vulnerabilities that he can exploit to gain more hours, ie low level sabotage that he is then paid to fix.	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	10 MINUTES	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	Paste Question HERE	
"The Event" - How will your person first hear about or experience the threat?	Other cities have adopted similar sanitation systems, management has announced the pending changes to union leadership, the city ran a pilot program to proof the concept. James understands once the system is fully online, he will be unemployed.	
Question Two	Paste Question HERE	
What will the person have to do to access people, services, technology and information they need?	James union and colleagues have a strong social network, sharing information on the system, furthermore social media allows them to learn from distant communities of interest (e.g. in other cities that have adopted the system). They learn and share techniques to disrupt the system by introducing false data.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Question One	Paste Question HERE	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	Government: automation has reduced each truck's crew from three to one, once the system is government certified, that policy will change to allow the city to terminate the last human operator.	
Question Two	Paste Question HERE	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?	In order to implement the automated situation, operators like James are required to ride along, monitoring performance and inputting data, i.e. training the automated system. However, they are now paid less per hour because the work is deemed less demanding and only work 50% of their former hours.	
PART FOUR- Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		

Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	City government must change existing regulations to allow automated sanitation.	City Gov.
2	James could re-educate himself to seek another job.	James
3	Globex Corp., vendor of the automated sanitation system, can monitor implementation looking for indications of malfeasance.	Industry
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
1	Loss of jobs as automation makes workers redundant.	James
2	Moral hazard to disrupt AI implementation to save jobs or attain other social aims.	
3	Threat of hackers exploiting vulnerabilities in AI system.	
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Research on implementation of automated systems that could displace large numbers of workers.	
2	Education and training programs to give workers new skills and career opportunities.	
3	Research social engineering that can disrupt AI systems.	
4	Research AI decision making.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Implement policies to take advantage of research results identified above.	

Group 5	
Experience Title:	Miguel the Migrant Worker - A Fruitful Future?
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	Machine self-adaption
Slot #2	Eventually you will take an 8 year degree to learn what happened in 1 month
Slot #3	Technology will be a means to enhance the human experience, not hinder it
Slot #4	AI will facilitate modeling/simulation of the medical field (specific cancers and drugs) through virtualization
Slot #5	Human mind can't keep up with data flow
Slot #6	By 2045 computers will have more intelligence than all humans combined
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Miguel is a 17 year-old migrant worker. He is first-generation in the United States, his dad died when he was young (raised by a single mother). He has two younger sisters aged 5 and 7. He is attempting to get a gymnastics scholarship at San Diego State. He's currently in high school doing well in Academics - extremely tech savvy. In order to help support his mom and two younger sisters, Miguel also washes cars as a side job. His mom works at a fruit packing plant and makes just enough money to support her and her three kids.
Where do they live?	Miguel and his family lives in Carlsbad, New Mexico. It was always his dream to travel in the United States and eventually settle down in southern California. Miguel loves the beach and temperate climate.
What is the threat?	Miguel is unsure of his future experience in college - he is focusing on the scholarship but isn't sure if his degree will be useful after graduation because of the increasingly rapid pace technology sets.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	Miguel's mom has the classic sense of the American Dream - she wants him to attend college, graduate with a degree in STEM, and move on into a highly-desired career field (he will be the first to attend college in the family). Miguel is analyzing this situation as an ambiguous future with too many unknowns. As a result, there is cognitive dissonance within the family - Miguel doesn't know where to go, and his mother wants him to support the family through the traditional career path.

PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	
Question One	"The Event" - How will your person first hear about or experience the threat?
	Miguel first attends a TED-type talk at his high school where a man presents the beginnings of AI. He states that the traditional college degree and career path will soon change. Miguel is interested in this topic, causing him to think about the impacts it may have on him and his family. This topic resonated with Miguel and made him excited. He believes that he wants to pursue this new path, but also wants to follow his mother's desires to support his family.
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?
	Miguel struggles to think about the current human strengths and how they will be out-classed in the near future. Miguel sees potential growth in certain career fields, specifically relating to computer engineering and software development. With this new development, Miguel realizes that there is a dichotomy and separation between the older and younger generations. The reliance on technology is growing, and now people in the United States struggle to keep up with the rapid pace of technology and growing artificial intelligence.
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	
Questions (pick two)	10 MINUTES
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
Business Models: What new business models and practices will be in place to enable the threat?	
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?	
Question One	Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?
	Technology and the continued progression of AI development will bring this dissonance about. AI and its repercussions will develop a cultural crisis or generational gap, similar to what is occurring today but on a larger scale.
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?
	The support necessary to bring about this threat are actors such as the military, government, and other companies that continue to develop AI capabilities. This is an inevitable threat that will occur at one point or another.

PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Development of education program focused on technology awareness for the older generation	Federal government
2	Organization of validation programs and third-party committees for AI career path development to society	DoD (because of current struggle for recruiting and retention)
3	School expansion of double majoring programs	Education facilities
4	Media and Hollywood "why we fight" for AI	Industry
5	Paid apprenticeship programs - internship idea for career development	Industry
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
1	Continued technology and AI development	All
2	Cultural acceptance of AI and similar capabilities	Society
3	Applicability of AI - how will the technology be employed in 10 years? (Medical? PTSD therapist use? Educational for online classes?)	Industry/Govt/Military
4	Use of AI is for malicious means/ends (ex. gang of con artists are using AI to get info from people - extension of a criminal enterprise) which would expand cognitive dissonance	Govt/Military
5	Disparate strategies for dealing with malicious AI (government and private sector have differing knowledge bases and differing goals)	Govt/Industry
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Education of older generation and ensuring information flow and validation of alternate career paths	
2	Advertisement and media outputs of artificial intelligence and documentaries of current and near-future capabilities	
3	Peace Corps beginning education and training younger and older generations - 'official and validated' path to success for American Dream	
4	Development of technology focused towards certain generations (think Jitterbug but with AI for older people)	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Education of older generation and ensuring information flow and validation of alternate career paths	
2	Grand nationwide challenge to highlight the future capabilities of AI and the possibilities for all generations that it may have	
3	Advertisement and media outputs of artificial intelligence and documentaries of current and near-future capabilities	
4	Peace Corps beginning education and training younger and older generations - 'official and validated' path to success for American Dream	

Group 67	
Experience Title:	Young millennial in a developing world
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	AI is its own cultural category(1)
Slot #2	AI could be more intelligent than Humans which could create catastrophic consequences/circumstance(5)
Slot #3	80/20 more important than ever(2)
Slot #4	Nations (or other groups) can take risks that individuals can't and vice versa(9)
Slot #5	Moore's Law- (As of 2017) human knowledge doubles every 12 months and the human brain cannot comprehend all of this information.(8)
Slot #6	Is the emphasis on STEM relevant?(20)
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Young person, still in school. 14 years old. Economically disadvantaged.
Where do they live?	Developing nation in Africa
What is the threat?	They need a career, a life. They need to train themselves to survive in the world, haven't had much training, and are losing job opportunities to AI
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	The adversary is their inability to find a meaningful life. They struggle to find a job with the spread of AI. We want them to avoid being trapped in their position on the fringes of society.
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	

Question One	What will the person have to do to access people, services, technology and information they need?	
	The nation will have to have incentives to provide educational services/training to their populous. Could we potentially make AI to deliver this technology and information cheaply and effectively.	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	
	If the necessary technology/AI isn't effectively distributed to this country, they may be stuck in the developmental stage for a longer time. This may end up rippling and attributing to the wealth inequality gap; also, if the students don't find a job and a way out, they may end up developing criminal habits. This may intensify conflict and put pressure on national security.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Business Models: What new business models and practices will be in place to enable the threat?	
	Capitalism drives innovation, and will be crucial in allowing the development of the country. Going forward, our business/economic models will require revision as automation and AI replace lower income/lower skill jobs.	
Question Two	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
	First, governments must want to enable their populations and provide them with the information/AI/technology to be successful and find jobs.	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Governments can control economic policy.	Government
2	R&D and innovation policy can drive development and sustainability.	Government/Private Sector
3	International aid, norm-development and international governance to decrease the divide between "haves" and "have-nots"	IGOs and NGOs
4	Education policy.	Government

5	Economic investments from industry	Industry
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Cultural limitations and politics; religion	Society
2	Geography/climate detirioration (we are past the point of no return)	Environment/Industry
3	Violent extremist ideology proliferation increasing because of the divide between haves/have nots.	Factions
4	Individual nations cannot control what the others do; the world remain anarchic.	Governments
5	Governments cannot control how individuals in developing countries perceive the world. As individuals become more connected, their cultures are harder to preserve.	Governments
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Innovation driven from outside traditional spheres	
2	Drive research funding for AI toward the countries who need the help and understand the problem. Also drive research toward social sciences, to include public policy, law, etc.	
3	Shape policy to improve education and mitigate impact of AI/automation on economically disadvantaged in developing countries. Delay the tipping point of a fully automated AI workforce by incentivizing companies to retain some degree of human talent or introduce automation tax.	
4	Identify who will lead efforts to disperse international aid for education, innovation, etc.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	cultural shaping of their vision of the future	

Group 8	
Experience Title:	Insider Threat to Process Control AI
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	Algorithms are not necessarily AI's
Slot #2	He doesn't see how AI DOESN'T destroy us - Path to AI taking over has already happened
Slot #3	80/20 more important than ever
Slot #4	You sometimes need to expose yourself to risk to gain credibility (WWI web of treaties - necessary to gain trust but with big consequences)
Slot #5	Technology will be a means to enhance the human experience, not hinder it
Slot #6	Decsion making must be performed at machine speed
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Mohamed Nakabale - Ugandan Process Control System Administrator
Where do they live?	Uganda
What is the threat?	Corporate competitor/organized crime entities that manipulate AI for Oil & Gas Process Control Networks to disrupt or degrade operations for financial advantage.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	Optimally, our person doesn't notice the threat is being experienced. The Threat Actor manipulates the AI to affect output levels in such a minute way as to disrupt output, but not enough that our person notices. This can allow organized crime or other actors to siphon off extra production, potentially to the black market. This is aided by the Black Box effect of AI--that our person doesn't really know what is going on. Threat actor wants to disrupt production and siphon off material. Our person wants to keep production in agreement with AI and be able to use the AI to do his job. Threat actor's actions can create distrust in our person for the AI. Insider threat could come from physical industry employees who have become subverted, spearfishing techniques, or modified AI updates.
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES

"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
	Delivered when our person discovers inconsistency with AI through random checks, what was asked for vs. what was delivered	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	
	People lose confidence in AI and algorithms as they will need to go to conventional processes while troubleshooting automated processes. Skill sets may not be immediately available to go back to the "future." AI was used to reduce overall operating costs. Increase in capital and operational expenses. Additional ripple effects include vendors having to rebaseline AI software and Ugandan forensics having to recover and rebuild.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	Weak governance structures prone to corruption which enable adversary to evade oversight. Corporate culture mirrors this, with conditions encouraging selling services to the highest bidder, rather than preserving integrity	
Question Two	New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
	Adversary uses extremist ideologies to influence inside contacts who assist in criminal behavior (i.e. modifying AI)	

PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Corporate quality control will account for more routine manual checks (better accountability)	Industry defender
2	Government institutes national infrastructure restrictions/controls	Government
3	Identify strict oversight of AI and people who can affect the AI, since the AI is critical to the industry performance	Industry defender
4		
5		
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Frequency and diligence of AI vendor software updates	AI Vendor
2	Predictable patterns in industry security, inspections,	Industry
3	Societal disruption that could allow industry employees to be compromised	Government
4	Supply and demand of oil and gas	
5		
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Support for strong financial institutions that can detect and prevent corruption	
2	Information campaigns to raise awareness of the threat of extremest organizations	
3	Implement strong insider threat awareness campaigns	
4	Identify interdependent systems which could be vulerable to attack	
5		
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Government oversight and regulations	
2	Out of band, unmanned, automated quality control checks	
3	Create public/private partnerships to improve Ugandan government structures to prevent corruption	
4	Institute controls in the AI to detect and repair external, unsigned modification.	
5	Invest in local education structures and exchange programs to enable local nationals	

Group 10	
Experience Title:	Do medical sheep dream in Hausa
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	"Artificial" carries connotations that are considered negative in today's society
Slot #2	What happens when the AI systems develop themselves and we the humans do not notice
Slot #3	We all leave a digital data trail, and that trail is going to be filled with richer and richer data as we deploy more sensors and more computing into the world. IOT, wearables, analytics, etc.
Slot #4	opportunity casting vs threatcasing
Slot #5	Virtual humans
Slot #6	STEM deficiency
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Harmony, an unmarried female Nigerian with an illegitimate newborn child that has been diagnosed with a significant genetic disease in an urban setting (Lagos) with a lower income. She works as a cashierist in a local business. Her family and hometown community has effectively ostracized her excepting for her sister and a few members, to include her employer.
Where do they live?	dense urban - Lagos, Nigeria . Nigeria is still a developing democracy with significant amounts of corruption within the government. There is still significant NGO influence within the country
What is the threat?	artificial intelligence doctors are trained off of a European dataset but then provided by an NGO sponsored startup to 3rd world communities.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	The newborn child has been incorrectly diagnosed with a genetic disease by an AI doctor. The mother can't afford a real doctor. However the AI doctor has been trained on an European dataset that does not take into account pollutants or other characteristics. We want them to be able to make an informed decision.
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	

When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	
Question One	When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?
	<p>The treatment makes symptoms worse instead of making things better.</p> <p>The mother will be worried and more and more frantic that her child is not getting better. She will also start trusting the AI doctor less and therefore start resorting to fake/ineffective medical treatments that are even worse.</p> <p>Later on an aid worker will realize that there is a locus of children dying in this area despite having access to the AI doctor. Upon further investigation they will realize that the urban area is polluted due to the e-waste recycling industry and that the AI doctor has been mis-diagnosing due to being trained on a european genetic/environmental model</p>
Question Two	How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)
	<p>Primarily electronic communciation (like an Amazon echo) that communicates the diagnosis and treatment recommendation. She submitted pictures/samples to the doctor electronically as well the day before.</p> <p>Every once in a while an aid worker will come around.</p> <p>Communication with government authorities is avoided due to corruption. Access to human doctors is limited to the very wealthy</p>
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	
Questions (pick two)	10 MINUTES
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
Business Models: What new business models and practices will be in place to enable the threat?	
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?	
Question One	Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?
	<p>AI doctors are trained with insufficient datasets because we have continued with today's anglosaxon focus in scientific research.</p> <p>This is marketed as the first solution to impoverished locations to provide medical treatment by an NGO sponsored startup. It was provided by the startup who is trying to help the city with all good intentions.</p>
Question Two	Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?
	The existing AI/ML methods and the existing medical research datasets will directly enable this threat to develop. Without intervention or changes in AI/ML training, this threat will occur

PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Require medical research against broader populations to produce broader datasets	Government
2	Invest in adversarial testing & training of ML/AI systems	Academia, Industry
3	Increase funding for training doctors in 3rd world countries to work along side the AI tools	Government, NGO, Universities
4	Increase benefits for doctors to move to 3rd world countries from 1st world countries (Increase funding to NGOs to increase access)	Government, NGO, Universities
5	Develop training program for physicians assistants to support AI doctors	Universities
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	success of AI doctors in 1st world countries	medical researchers
2	New massive VC investment in AI/ML doctors for 3rd world countries	VC, Entrepreneurs
3	Less than .5 doctors per thousand people in Nigeria	Nigerian government
4	AI doctors become more cost effective than providing real doctors to a region	entrepreneurs
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	educate UN and AU policymakers on the importance of diverse sample in medical research before medical treatment is distributed to impoverished countries	
2	Fund research into adversarial/attack activities against AI (ala https://blog.openai.com/adversarial-example-research/)	
3	Evangelize the need for more human doctors/PAs who know how to work along side AI. Do this with to NGOs, Philanthropists,	
4	Define PA for AI scope of training, popularize the concept	
5	Improved STEM education at the younger grade levels so they are prepared to obtain a P.A or M.D	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	PA for AI program needs to be developed and implemented and people need to be convinced to start pursuing this career	
2	Make adversarial testing of all AI products an industry standard BKM	
3	Set policy around requiring diverse populations in medical research in order to achieve government/NGO funding	
4	FDA (or other approving body for medical AI systems) sets new requirements for showing diversity in training data used for all AI products	

Group 11	
Experience Title:	Sleepless in New Jersey- Destination, Port Fear
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	AI is another manifestation of what means to be human
Slot #2	Artificial Intelligence will destroy us
Slot #3	Cleared individuals are at more risk than ever and it's getting worse.
Slot #4	Focusing on negative creates the possibility to ignore the positive
Slot #5	By 2045 computers will have more intelligence than all humans combined
Slot #6	In 10 years, the width of transistors will be 5 atoms... Moore's Law will fail according to physics
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Native American Male, 45, former longshoreman, single father, GS-11 Lead Customs Inspection Specialist at CBP assigned to the Port Authority of NY/NJ
Where do they live?	Clifden, New Jersey, USA
What is the threat?	Microtargeted by foreign intelligence using big data/AI on social media, etc. to use his personal life to interfere with his professional responsibility. For example, changing a negative on a mandatory drug test to a positive. Or falsifying timesheets to put him under disciplinary proceedings. Ultimately to take him out of rotation. Use cascading results from targeted individual to introduce a compromised shipping crate containing a dirty bomb or a large EMP generating device to kinetically assault the port infrastructure.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	Professional Disciplinary actions predicated upon false data constructed by a potential state bad actor.
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	

What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	What new capabilities enable the person and their broader community to recover from the threat?	
	Redundant data systems that can disprove or falsify weaponized data. We don't want to rely on a single source to attest to an individual's activities.	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	
	Once a weak human link in the chain is eliminated, an adversary has an entry point into the physical system. For example, smuggling a dirty bomb or EMP into now-uninspected cargo at this giant port. Even short of this, compromising this individual also makes it that much easier to compromise others (such as supervisors) or even recruit the now-compromised individual as an informant/intelligence asset.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Business Models: What new business models and practices will be in place to enable the threat?	
	Social media and "digital exhaust" related to outside-world practice that might be attractive targets for compromise. E.g., travel/gambling/drinking habits, questionable social media posts, and so on.	
Question Two	Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
	Proprietary algorithms and "big data" approaches that can sift through the massive amount of data (above) on millions of social media users and integrate them to pick out "ripe" targets for micro-targeting.	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Access to employee data (with agreement with employee)	Employee and government
2	Government computer systems (to detect intrusion)	
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?

1	Proliferation of personal data in the digital realm	Private companies
2	Profligate use of social media, etc. by employees (e.g. putting irresponsible material on them)	The user
3	Advancements in algorithms/etc. that allow adversaires to make more hay out of more data	
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	New relationships between government and social media, medical records keepers, etc. that ensure standards of security (both for the user's benefit and the employer, in this case the government).	
2	Carrots and sticks: shields from legal liability for hacking IN EXCHANGE for adequate standards.	
3	Educate the government workforce on the existence of and inherent risks involved in our "everyday online life." Populate the concept of "digital exhaust" and how related PII and activities can be exploited as part of larger risks beyond the targeted individual to the organization of which they are a part.	
4	Allocate resources within the ODNI and the broader IC on these risks to construct statistical models on infrastructure and personnel to determine where the compromising of "weak links" could exact the most damage and take actions to educate and prepare action plans.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Complementary technology to simple passwords (e.g., embedded nano-technology) that can certify that the actual user is the one who updated her social media or accessed sensitive records.	
2	Predictive Analysis based on your digital profile and the aggregation of many data sources "Minority Report-like, pre-crime" in the context of vulnerability. What is your risk-factor/risk-score	
3	AI scenario modeling of the implications, cost/benefit analysis of mitigative actions based on the data in 2. I.e., is it better to just let go/not hire a "weak link" identified in 2 or does that cause more problems than it solves?	

Team: 12	
Experience Title:	Let's Dance
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	Humans have long told stories and myths about animation & bringing things to life.
Slot #2	We do not not know how long to create Super Intel AI safely
Slot #3	OSMINT and SOCMINT (social media intelligence) will matter more in the future than they do now. More information will be available on people as they spend more of their lives online in social media platforms. Social media will evolve rapidly to add virtual/augmented reality, voice and gesture. Connections will continue to expand beyond "friend-to-friend" to connect people to services, markets (a full-fledged transactional platform), businesses, and many other organizations.
Slot #4	New markets occur all the time
Slot #5	Human mind can't keep up with data flow
Slot #6	Will we need "Programmer Archaeologists" to help modify SW code to keep pace with AI?
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Ba Wei is a young female undergraduate studying computer science, trying to figure out the world. She's from China, from an upper-echelon state functionary family that lived in a city, and has migrated away to get away from pollution. She now lives in the United States in a Chinese diaspora.
Where do they live?	She lives in Portland, Oregon.
What is the threat?	The girl is exploring new markets, (social and economic markets) trying to determine where and how she can fit in. She is torn between loyalty to China and her family versus loyalty to her new life in the US. Her parents are trying to maintain control over her life through monitoring her social networks and marketing information. Her parents sense a threat from a larger force (likely from China) that is monitoring all of their lives to ensure adherence to nationalistic norms.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve?	
What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	

	<p>The Chinese government is using a super-capable AI to monitor the girl's family, and use predictive analysis to nudge her back on the proper course. She does not know the actual process the AI is using to monitor her activity, but she sees its injections into her everyday life. She sees suggested ads, programs, etc. that seem to be in line with Chinese nationalistic norms, but don't fit with her normal social interactions and internet use. We want the girl to live a life unencumbered and unfettered by the Chinese AI and for her to pursue more liberal economic and social interactions not controlled by China - to be more of a citizen of the world. We want to avoid the AI being used as a surveilling force, and instead as an enabling force. We want her to avoid the disintegration of her family due to competing values and loyalties.</p>	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	10 MINUTES	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	"The Event" - How will your person first hear about or experience the threat?	
	<p>The girl was working on a project in one of her classes and analyzing the use of artificial intelligence for mass monitoring. She discovered indicators that her experiences with suggested ads, programs, etc. were actually from the Chinese government's AI. She realized that her parents were complicit and attempting to conform her behavior. Her family was receiving information from the AI, and actually acting as agents of the AI.</p>	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	
	<p>The broader implication is the replication of this technique by other nations - this is not an aberration, it's normative of all societies. There are the implications on individual liberty, economic freedom, and self-actualization. There are international relations implications as well. Her family's collusion with the AI has implications for the human-AI relationship.</p>	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	

	There are multiple threats to controlling the girl's loyalties. At the state level, the United States will not knowingly allow influence activities against US citizens. The threat of open-source intelligence, (how the girl found out about the AI in the first place) has implications to the Chinese trying to use the AI. There are also corporate and economic barriers to what amounts to hacking into someone's social and internet life. There may be a range of legal and illegal counters to the AI trying to control Chinese citizens abroad - some public-sector, some private-sector.	
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	The Chinese will need to acquire or create companies that can produce or control AI platforms. They'll need industry partners or industrial knowledge that will get them access to operating systems, internet companies, software developers, etc. They will want or try to gain approval or a blind eye from other states - they will need to have strong diplomatic soft power to mitigate international backlash. They will need access to whatever amounts to a public ID in 2027 - a device under the skin, a data tattoo, etc.	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Technical acumen, cultural access, hacking ability	Chinese hackers
2	Access to knowledge of technology at/near state-level	The girl / international hackers
3	Tools and the will to use them against state-sponsored AI	The girl / international hackers
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
1		
2		
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1		
2		
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1		
2		

Group 14	
Experience Title:	AI Insurrection of Benevolent Actors
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	Inventors of AI were reacting to/against Freud & messiness of EU psychology
Slot #2	Culture matters when interpreting observations
Slot #3	Algorithms are not necessarily AI's
Slot #4	AI is its own cultural category
Slot #5	Humans have long told stories and myths about animation & bringing things to life.
Slot #6	AI is actually a broad constellation of technologies -- many of which are extremely dissimilar from each other
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Ziva David, female, had to flee home village that was overrun by rebels.
Where do they live?	Refugee camp in Somalia
What is the threat?	Autonomous systems assign resources(food/water/education/medicine) that threaten Ziva's survival
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	The camps resources are managed by a Western European organization. This organization has attempted to leverage the latest algorithms from silicon valley. However, the algorithms that optimize distribution of resources are corrupted by a 3rd party who wants to maintain instability in the region. The aid organization, technological hubris, defers all decisions to the AI system. Increasingly, the camp inhabitants don't understand the distribution decisions of the resources as it's not immediately logical. The 3rd party aims to sow insurrection/instability to consolidate power. Ziva's objective is some agency in her survival and movement to another nation, the threat actor is incentivized to keep camp inhabitants fighting amongst itself and the aid group. Ziva would like to not experience malnutrition and under-education
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	

What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	"The Event" - How will your person first hear about or experience the threat?	
	Ziva does not know directly of the threat actor, and sees the poor distribution of resources as due to mismanagement and trust in resource allocation algorithms. Rumors abound in the camp about corporations performing societal experiments, purposeful neglect, or lack of money and Ziva cannot connect her lack of resources to the threat actor	
Question Two	What new capabilities enable the person and their broader community to recover from the threat?	
	Ziva confers with her members of her group and they decide to completely eschew the aid organizations distribution decisions and rely on conventional human networks to communicate needs and hand out resources.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
	Crowdsourcing is used to perform R&D of the algorithms and artificially induce resource disparities	
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	The democratization of computer science skills reduces the R&D infrastructure needed to pursue "software" weapons	
PART FOUR- Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Conventional 'systems' that don't rely on optimized algorithms(p2p communication)	Government/NGO
2	Detection of anomalies in the environment(both inputs and outputs) that would alter AI abilities to efficiently distribute resources	Government/NGO

Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Due to global education inequality, the less educated will not have any defenses to disrupt/mitigate cyber effects where previous physical combat did not have minimal education 'requirements'	Everyone
2	NGO reliance on AI to efficiently distribute resources, if compromised, erodes refugee trust in any formal organization that AI in any form	Threat actor that compromises what was perceived as purely altruistic endeavors
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Networks of expertise for cyber-intrusion, cross-agency coordination	
2	Crowdsource defense of systems.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1		
2		

Group 15	
Experience Title:	AVOID JUST BECOMING A NUMBER,
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Slot #1	"Artificial" carries connotations that are considered negative in today's society
Slot #2	Humans have long told stories and myths about animation & bringing things to life.
Slot #3	Convenience in cyber applicaiton can lead to inadvertant shortcuts
Slot #4	opportunity casting vs threatcasing
Slot #5	By 2045 computers will have more intelligence than all humans combined
Slot #6	In 10 years, the width of transistors will be 5 atoms... Moore's Law will fail according to physics
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Female fast food worker, 20 years old
Where do they live?	big city
What is the threat?	She has experienced co-workers losing their jobs and the community suffer from economic decisions on the part of the corporation, as well as decrease in product quality
Briefly describe how your person experiences the threat. SHE IS CONCERNED ABOUT LOSING HER JOB TO AUTOMATED SYSTEMS AND HOW THE COMMUNITY IS AFFECTED BY LOSING NEIGHBORHOOD SERVICES... HOW MIGHT AI HAVE NEGATIVELY AFFECTED FOOD SAFETY/QUALITY BY TRYING TO MAXIMIXE PROFIT VS QUALTIY. HER JOB MAY BE IN JEOPARDY IN AN EFFORT BY AI TO IMPROVE EFFICIENCY BY ELIMNATING HUMAN DECISION POINTS	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of? SHE IS CONCERNED FOR HERSELF/FAMILY AS WELL AS COMMUNITY FACTORS. THE THREAT ACTOR IS CONCERNED ABOUT THE COMPETITION, CUTTING COSTS, GAINING RELEVANCY, AND KEEPING PACE WITH CHANGE. PROFITABILITY.	
What is the experience we want the person to have with the threat? ACHIEVE SUCCESS IN PROVIDING FOR HER FAMILY. SEE OPPORTUNITY TO ENHANCE AI APPLICATION BY INCLUDING HUMAN FACTORS IN THE PROCESS	
What is the experience we want them to avoid? AVOID JUST BECOMING A NUMBER, NEEDS TO BE RELEVANT. RELY TOO HEAVILY ON AI MAKING DECISIONS BASED SOLELY ON A GIVEN SET OF DATA	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	10 MINUTES
"The Event" - How will your person first hear about or experience the threat?	She sees compromise in qualtiy and services to the community

What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	She is the interface between the customers and the product and sees firsthand the reaction to corporate changes in food quality and service	
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	"The Event" - How will your person first hear about or experience the threat?	
	She may hear about job cuts in her industry on the news or social media. She sees compromise in quality and services to the community and fears loss of her job. She is the human interface between the product/services and the consumer and hears both from her employer and her circle of family, friends, and customers	
	She elevates her food safety concerns to her supervisor and/or public media;	
Question Two	How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
	She is the interface between the customers and the product and sees firsthand the reaction to corporate changes in food quality and service	
	She will be given direction by her employer, but will provide feedback to family and friends informally, social media, personal contact	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	10 MINUTES	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
	Acceptance of the choice between jobs vs. commodity; food quality/safety vs. profit	
Question Two	New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
	Emerging technology/AI will help bring more affordable options; and the food industry will strive to maximize profit	
PART FOUR- Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		

Gates:		
What are the Gates?	things we have control over...	
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	food handling regulations; inspections;	local/state/federal govt
2	improve research on food safety and production efficiencies	federal govt, industry, academia
3	occupational counseling	industry, govt
4		
5		
Flags:		
What are the Flags?	things we do not control...	
	do not control knowledge...	
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
1	cannot restrict access to information/knowledge	govt?
2	cannot restrict corporate greed to maximize profit	industry
3	cannot influence natural disasters that may impact food production	N/A
4		
5		
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1	identify food handling protocols	govt
2	inform public on pros/cons and alternative occupations to compensate for technological advances	govt
3	anticipate threats to food safety and put mechanisms in place to address vulnerabilities	govt, industry
4	research standards between govt and industry	academia, industry
5	subsidize positions/salary of critical positions	govt
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	invest in R&D to identify technologic advances that could limit impact	
2	basic, applied, and advanced prototype development investment to enhance food safety and security	
3	agreement on standards between govt and industry	

Group 1	
Experience Title:	The Perfect Red Team
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	State Sponsored
Slot #1	AI is actually a broad constellation of technologies -- many of which are extremely dissimilar from each other
Slot #2	He doesn't see how AI DOESN'T destroy us - Path to AI taking over has already happened
Slot #3	Cleared individuals are at more risk than ever and it's getting worse.
Slot #4	New markets occur all the time
Slot #5	Virtual humans
Slot #6	Collapse of Moore's Law
Wild Card	Supply chains can be controlled by AI
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	MAJ Don, mid-level expert in an auditing/evaluating AI-controlled supply chain. He oversees a team of auditors for the Army's AI-controlled supply chain.
Where do they live?	Hill AFB, Ogden, Utah
What is the threat?	AI system is compromised and is controlled by Russia.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of? The AI in the supply chain is responsible for monitor and managing readiness status of vehicles/weapons systems. It can track the information on training and actual operations.	
What is the experience we want the person to have with the threat? MAJ Don discovers inconsistencies with MRE packaging for humanitarian aid deliveries in Maldives. Based on this inconsistency, he investigates the security of the supply chain and uncovers other minor inconsistencies in warfighter supplies and realizes that the AI system has been compromised. DoD is planning a multi-force engagement in Syria to begin in 13 days.	
What is the experience we want them to avoid? In 2027, DoD supply chain is completely controlled by AI that has created inconsistencies in food supplies for humanitarian assistance. Russia has introduced a worm into the AI. Wants to erode trust/create uncertainty in AI and introduce confusion across supply chain and operations, motivating US DoD to revert to human control of supply chain. This action would temporarily disrupt DoD supply chain, allowing Russia to exercise strategic control of pending action in theatre in Syria. Discover the threat and disrupt the threat Avoid notifying the AI system that the worm/inconsistencies have been discovered. We don't want him to cause panic and disrupt the supply chain midoperation	

PAUSE : Call in a facilitator to discuss / debate before moving on	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	15 minutes
"The Event" - How will your person first hear about or experience the threat?	He notices the order of oral rehydration salts for US PACOM quadrupled due to a diarrhea outbreak from spoiled MRE's with mislabelled expiration dates. This raises a red flag for him because MRE's usually lasts for 10 yerars.
What is different and/or the same as previous events or instantiations of the threat?	First time to see an anomaly since it has been an AI-controlled supply chain. The replacement of humans was made to prevent the mislabelling mistakes
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	He will not realize until later that the AI system has been compromised for more than ten years. He discovers that Russia has been gathering information on military operations using the supply chain AI and has been learning the behavior of the organization in an attempt to build an Virtual Red Team (using a different AI systepm) representing the US. The worm can send location/status of personnel and equipment during all the training exercises and the AI constructs an image of US strategic and operational decision making using all the information it gathered from supply chain movements during peacetime. Decision made to extend the expiration date for MREs by 4 years - relabeling, but not communicated.
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	Aside from gathering information on troop/equipment movement around the world, it is also possible that a dormant worm can be gathering information on the personality/decision-making behavior of key leaders and it allows them to predict their response to all possible war scenarios. They can create a virtual image of the mind of individual leaders.
Question One	Paste Question HERE
	Answer Question HERE
Question Two	Paste Question HERE
	Answer Question HERE
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	
Questions (pick two)	15 minutes
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	Automation of the supply chain/ equipment maintenance/upkeep system is much desired in the DOD because it a demands less DOD support personnel
Business Models: What new business models and practices will be in place to enable the threat?	
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	Digital tagging of all supplies/equipment to enable full automation of the supply chain system.

Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Paste Question HERE	
	Answer Question HERE	
Question Two	Paste Question HERE	
	Answer Question HERE	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		
1	Decision to fully automate	Who is the responsible party? DOD
2	Selection of contractors/suppliers	Civilian defense contractors
3	Funding of AI supply chain development	Congress
4	Reporting/auditing practices for AI systems	Accounting/auditing professional organizations
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearldis of the future to come.		
1		Who is the responsible party?
2	Quality of residual human actors in the supply chain system	Individuals
3	Societal expectations/acceptance for increased automation	Society
4	Societal expectations for general accounting practices	Society / Professional organizations
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Secure transition to AI systems including redundancies and contingencies to go back to manual supply chain auditing/evaluation	
2	Train workforce to manage transition from human to AI systems	
3	Create standards and tranparent practoes for AI system auditing including mechanisms to ensure that systems are not introduced with problems or compromises inherent	
4	Plan to engage and transition supply chain workers from human to AI system	
5	Fully automate systems that will be under AI control	Delineate which segments of supply chain will remain under human control
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		

1	Train workforce for transition from AI to human systems
2	Increase and diversify the auditing workforce team to ensure adequate human oversight over continually-learning AI systems
3	DOTMLTF-C must be structured in a way that we can resolve auditing inconsistencies without threat to the intelligence security of the system.
4	AI security auditing practices must be developed (similar to cybersecurity auditing practices that must presently be developed for medical devices)
5	Develop AI capacity to self-regulate, self-diagnose anomalies.

Group 213	
Experience Title:	D-Day 2: Port Unauthorized (starring Kevin Bacon as Mayor Bacon and Rob Schneider as a bridge named "Deuce Bridgelow")
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	State Sponsored
Slot #1	7- AI reflects the worldview/biases of its creators
Slot #2	6 - Artificial Intelligence will destroy us
Slot #3	2 - 80/20 more important than ever
Slot #4	4 - power of coalitions
Slot #5	10 - AI replacing work
Slot #6	5 - Collapse of Moore's Law
Wild Card	Slot 2 -There is an unregulated race to create the first Super Intel AI
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Mayor of Galveston, community is shipping, energy industries, medical school
Where do they live?	Galveston, TX

What is the threat?	Mayor wants to integrate AI into the Galveston port authorities to boost the economy. The mayor pushes legislation to roll out AI into the port as soon as possible so that there can be a noticeable improvement in the local economy before the next election cycle. The AI software company that won the bid for the contract was a local startup looking to take advantage of the new market. In order to minimize cost, the startup company buys one of the AI parts from a Russian supplier. This Russian supplier is sponsored by the Kremlin in order to manufacture these parts. The procurement process did not vet the company properly before giving the order to integrate AI into the port.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	Foreign government access into local industry systems. An entire shipment of contaminated food passes through the system and is moved through the Southwest. People start dying and they trace the cause back to the food shipment in Galveston. The mayor realizes that the AI system is not working as it is supposed to and that there are shipments that have been coming in for the last 18 months that have been shipped to various parts of the US that they cannot track down. They shut down the system and enlist the help of federal authorities for disaster relief efforts.
PAUSE : Call in a facilitator to discuss / debate before moving on	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	15 minutes
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	
Question One	"The Event" - How will your person first hear about or experience the threat?
	It will be in the news before he gets a report about the security flaw. The effect is known before the cause and location are identified. It is quickly discovered that the security of the system has been compromised which causes protests and civil unrest in Galveston, completely disrupting the local economy. The public loses faith in AI systems as a means of augmenting industry profits and shipping in Galveston comes to a complete halt.
Question Two	When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?
	The mayor will not see any threat to begin with. If anything, they will see a boost in the local economy due to the AI implementation. After deaths start occurring and the feds approach the mayor about the situation. The mayor will downplay the issue to the public to calm their fears while. What they don't understand initially is that there have been shipments throughout the entire US to unknown locations (that were modified by the AI). The NAional Guard collapses in on Galveston, believing that it is an isolated incident, before recognizing the bigger issue across the US.
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	

Questions (pick two)		15 minutes	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?			
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?			
Business Models: What new business models and practices will be in place to enable the threat?			
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?			
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?			
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?			
Question One	Business Models: What new business models and practices will be in place to enable the threat?		
	outdated acquisition process, underpaid procurement workers		
Question Two	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
	Russia need to identify the technology and maintenance of the systems in place,		
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)			
Examine the combination of both the Experience Questions as well as the Enabling Questions.			
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.			
Gates:			
What are the Gates?			
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.			Who is the responsible party?
1	Updating the acquisition process		procurement agents, state government cyber QC
2	Redundancy and cooperation between human and AI actors in the shipping process. Using AI as an augmentation of security rather than total replacement		port authority
3	Creating security standards for artificial intelligence		Academic, Silicon Valley, Private and public partnership
4	Adapting the process to check the security of AI systems over the next decade		Federal agency
Flags:			
What are the Flags?			
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as harbingers of the future to come.			Who is the responsible party?
1	Proliferation of AI systems. Two industries may never talk to each other		federal government establishes regulatory guidance
2	Private industry procurement		private companies
3	Geopolitical threats		
Milestones:			
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.			
1	Revisit state and federal acquisition guidance		
2	Develop public and private partnerships for AI standards		
3	Promoting utilization of AI systems as a means of security augmentation rather than complete job replacement		
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.			
1	Actual legislation that integrates AI into Homeland security		
2	AI software development that checks other AI software (FBAI?)		

Group 93	
Experience Title:	1984 in 2027: Our Orwellian future of transparency and AI
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	Criminal Organization
Slot #1	Humans have long told stories and myths about animation & bringing things to life.
Slot #2	Humans are incapable of establishing appropriate responses to AI
Slot #3	Convenience in cyber applicaiton can lead to inadvertant shortcuts
Slot #4	New markets occur all the time
Slot #5	AI replacing work
Slot #6	Will we continue to be able to conduct covert and/or clandestine operations in the future? (Given that sensors will be everywhere and transparency will become dominant)
Wild Card	Augmented reality is the primary computing platform for most people
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Harriet is a 38-year old programmer of AI bot trading algorithms for Goldman Sachs. Works in London in The City. Commutes by train each day. She's married with two young children (Pendragon and Nigel).
Where do they live?	Sevenoaks, UK
What is the threat?	Blackmailed with footage of an indiscretion hacked from her augmented reality glasses, Harriet is coerced to change the algorithms behind Goldman Sach's AI bot trading platform to favor a Russian criminal gang. Video streaming from her glasses was hacked from Apple's video servers. The criminal gang found her by using an AI bot to scrape social media (to learn where she worked), monitoring her movements by analyzing her phone GPS, and a wide range of public CCTV/Wifi data. They also hack her private email accounts to monitor them for anything they could use against her.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	

	Experience we want her to have: Harriet is contacted by Goldman Sachs IT saying that their Risk AI estimates she is at a 83% chance of blackmail and her access to IT assets has been frozen pending a further investigation. She is offered support to defend against any pending threat and given resources to mitigate it. Experience to avoid: Harriet is contacted by the criminal gang who replay a video recording of her indiscretion to her through her augmented reality glasses. They make a clear demand that she insert rogue code into the GS AI bot and threaten to show the footage to her husband if she doesn't comply.	
PAUSE : Call in a facilitator to discuss / debate before moving on		
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	15 minutes	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	What is different and/or the same as previous events or instantiations of the threat?	
	Harriet is contacted by the Russian gang, presented with compromising video, and blackmailed. She provides the gang with information on the AI bot algorithm. In a couple of weeks she is given a piece of code to inject into the bot that causes Goldman Sachs to sell millions of shares of several stocks that the gang has shorted in the market. Just after she has inserted the code, she is approached by the Goldman Sachs risk assessment team. Their risk AI bot was monitoring her activities and raised a flag when it noted changes in some of her patterns...specifically that she and a co-worker have been spending time together outside of work hours. They sensed this by scraping city CCTV and using face recognition to spot the two lovers in a variety of locations around the city. She has therefore been assessed as being at high risk of blackmail. They shut down access for Harriet, but it's too late.	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	

	<p>The accuracy of the AI risk bot is only as good as the quality and volume of data that it can access. The data such a bot might access includes work email, work systems access, and public social media feeds. Such a bot would work much better if it started to invade the privacy of employees. For example, the bot might scrape more private information: private email accounts, messaging services, GPS data from phones (which may be company-owned assets), video streams from augmented reality glasses, video streams from cameras located in the office, video streams from the hundreds of thousands of closed-circuit TVs located throughout the city of London, personal credit card data, and by reading MAC address from WiFi hotspots etc. The bigger implication of all this: How much privacy will be invaded to feed the Risk AI bot? How much violation of privacy will employees accept? Will they accept more if they hold a sensitive position with access to IP, money or other assets? Since every public space is now transparent, anyone can run analytics to understand more about who is going where, who with, and what they are doing, especially as video analytics become more sophisticated and lower cost. Work environments will also become transparent as companies deploy more sensors in the work space.</p>	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	15 minutes	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
	<p>TODAY: Data mining, GPS tracking, MAC address sniffers (like the ones in the trash cans in London...), cameras, facial recognition algorithms, voice recognition (and voice-print mapping), behavioral analysis tools, code injection techniques (cyber hacking kill chain). FUTURE: 1/ AI algorithms to identify optimal targets based on their access to assets, and their likely vulnerabilities. 2/ Tools and a dark web market for a service that is able to identify the exact target needed to manipulate a particular stock, set of financial assets, or that has access to a particular piece of IP.</p>	
Question Two	Business Models: What new business models and practices will be in place to enable the threat?	
	A dark organization that offers the ability to manipulate stocks or other financial assets on demand. They use AI to find and compromise targets with the appropriate access. The AI identifies the individuals that have the needed access. The AI also finds their vulnerabilities and makes specific recommendations on the best strategies to exploit them (AI-supported prescriptive analytics approach).	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		

Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Protective AIs - Both offensive and defensive	Defensive (Goldman Sachs) Offensive (Government Response)
2	Public Data Collection Points [Cameras, Hot spots]	Local government, industry
3	Security capability to track individuals by analyzing publicly available CCTV and WiFi hotspot data and running analytics to cross-reference face recognition data and MAC-sniffing algorithms to identify MAC address and track.	Industry, government, military
4	Access control systems tied to Risk assessment AI	Industry
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		
1	Under pressure from citizens who are concerned by the escalating rates of terrorist attacks on their cities, local governments move to widely deploy CCTV cameras and MAC address sniffers in major cities all over the world.	Local city government, industry
2	Responding to industry demands, the major cloud computing vendors deploy AI capabilities available on-demand in the cloud. This lowers the barriers of entry for individuals wanting to use AI for nefarious purposes and changes the cost equation for using offensive AI to identify and attack targets.	Industry
3	Weaponization of AI for illegal activities; using technology for other than intended	
4	Passwords are increasingly replaced with biometric access to increase security levels, and/or improve efficiency. Biometric collection eases access to content and other capabilities (such as work assets, payment mechanism etc) but means that databases are now being built filled with huge amounts of hackable biometric data (face prints, finger prints, voice prints, iris scans etc)	Industry
5	The increasing complexity of AI systems and code make mitigation and recovery very difficult.	Industry and Government
6	Wide range of voice-enabled devices and appliances are rolled out in millions of homes (all fitted with microphones, and some with cameras). These devices, because they are very cost-focused, have very low security, or no security features at all. Examples might include connected toys (for example, Cognitoys dinosaur), washing machines, as well as the voice-assistant devices like Echo.	Industry
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1	Greater implementation of encryption (especially end-to-end). Increased user privacy will make it harder for organizations to access the data in the first place (for exploitation).	
2	Improvement in predictive analytics techniques and insider threat capabilities. Identifying individuals who are susceptible allows organizations to approach them before an issue occurs.	
3	Increased public / private coordination of cybersecurity efforts. Sharing data will enable organizations to perform more informed counter exploitation operations.	
4	Increased layer 2 obfuscation implementation (obfuscating mac addresses) by default for end user. This will make it difficult for individuals to be tracked by nefarious organization (regardless of intent).	
5	AI capability for hunt and identification of vulnerabilities	

What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Maturity of AI algorithms and threat-based analysis. As AIs grow, protective capabilities will enhance as well as opportunities for exploitation.	
2	Increased privacy legislation passed. Placing an increased value on privacy and a higher level of protection will inherently make collection of personally identifiable information more difficult.	
3	Turn on Einstein (Intrusion detection capability). Nation-wide analysis of traffic and development and recognition of nefarious signatures and behavioral characteristics will make it easier to identify when something malicious is occurring.	
4	Hardware components all have security designed-in as the default for all IoT devices. Less dumb smart things, more smart smarthings. Not all of the device's intelligence is in the cloud, but also on site.	
5	Consumers will be buying smart, connected devices by the billions. From toys to appliances and everything in between. They will begin to demand an easy way to see if an item is "secure" during the purchase process, much like the Energy Star logo for energy-efficiency, or the CE mark for safety. This security compliance certification would cover consumer and commercial devices. An organization (which could come out of an existing certification organization like CE) would do dedicated testing and analysis on devices. They would maintain up-to-date standards and a methodology that organizations must follow to keep their devices secure (including end-to-end physical security, authentication, software update mechanisms, etc).	

Group 5	
Experience Title:	From China with Love
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	Criminal Organization
Slot #1	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.
Slot #2	Virtual humans
Slot #3	Machine self-adaptation
Slot #4	AI replacing work
Slot #5	By 2045 computers will have more intelligence than all humans combined
Slot #6	AI will facilitate modeling/simulation of the medical field (specific cancers and drugs) through virtualization
Wild Card	United States targeted by numerous terrorist organizations
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Zhang Wei is a 55 year-old PLA Commander who has many questionable relationships within the Chinese military. Over the course of his career, Wei has had ties to organized crime units in China, terrorist organizational ties, and a number of bank scandals relations. However, through his amassed power from these relationships, Wei has continued a successful within the PLA. In order to successfully smuggle in the dirty bombs to the US, Wei partnered with a third party hacking organization in China that developed an AI capable of creating a large network of virtual humans. This allowed for the US police, coast guard, and customs personnel to be distracted by virtualized humans rather than focus on the real threat of dirty bombs smuggled in. Additionally, the bombs were smuggled in to the US via the hacking of the logistical chain present within newly developed smart cities.
Where do they live?	Zhang Wei lives in China, and travels throughout the world for multiple meetings, conferences, and events as a representative of the PLA

What is the threat?	ISIS/ISIL as a terrorist organization became an established caliphate within the Middle East. Many nations actively oppose them both militarily and publicly as a legitimate nation. As a result, ISIS/ISIL managed to take hold of an unknown number of nuclear weapons that were unaccounted for by Russia following the end of the Cold War. All weapons were smuggled into the US, four of which were placed in major cities (Seattle, Houston, LA, and Chicago). All were detonated simultaneously within the cities, creating mass chaos within the US. The terrorist organization stated that they have a number of additional dirty bombs already within the US, and they will detonate the bombs if the US does not follow a series of demands - mostly monetarily on part of the terrorist organization. As a result of these attack on the US, China has taken the initiative to mobilize forces in the South China Sea, taking over the Spratly, Paracel, and Senkaku Islands. The US, focused on domestic affairs, is unable to counter this rapid mobilization of Chinese troops in the region. OVERALL: The threat that occurs is China's plausible implication in the attack on the US - the US suspects that someone funded and backed ISIS/ISIL and believes that another player is behind it (likely actors: Russia, China, Iran).	
Briefly describe how your person experiences the threat.		
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?		
What is the experience we want the person to have with the threat?		
What is the experience we want them to avoid?		
	Zhang Wei is in China on duty while this event occurs within the US, and he soon realizes that if he is exposed, his entire family, himself, and his colleagues will be at risk. Because of this, Zhang Wei must do everything he can to avoid exposure.	
PAUSE : Call in a facilitator to discuss / debate before moving on		
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	15 minutes	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	"The Event" - How will your person first hear about or experience the threat?	
	Zhang Wei first experienced this threat by hearing about the detonation of dirty bombs on the news. China and the rest of the world (aside from Zhang, ISIS/ISIL, and a few Chinese officials) were extremely shocked at the occurrences. Following the event, Zhang suggests to the head of Chinese forces that this would be an opportunity for China to expand on island-building and force mobilization in the East and South China Seas.	
Question Two	What will the person have to do to access people, services, technology and information they need?	
	Zhang Wei will require the logistical and support chain through his number of questionable relationships over the course of his career. His criminal support will direct him to the third-party hackers for the AI bot development, and his past operations will direct him to the ties with ISIS/ISIL - all of this allowing for his plan to infiltrate the United States and ISIS/ISIL.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	15 minutes	

Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
	Development of smart cities with revolutionized logistical chain, AI expansion allowing for continued bot development that generate fake people (unknown to all except the program), and dirty bombs following the Cold War era are located on the black market; Zhang Wei will utilize a number of different criminal organizations in order to create a complex network with multiple chains in order to create as much confusion as possible with regard to the responsibility of the actor	
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	ISIS/ISIL terrorist group has a will to execute a mass chaos attack on the US; AI software capable of developing in-depth human bots with realistic traits (i.e. virtual humans with relationships spanning the course of years); adaptive self-learning and intuitive online entity able to infiltrate a nation's networks; ability to find ways around and plan for the counteraction of defensive systems and networks (i.e. zero day virus on massive scale with loopholes); AI capable of developing layers upon layers of false trails based within real systems and attached to multiple enterprises (both criminal and not criminal)	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Initial proxy (ISIS/ISIL) used as a distraction	Zhang Wei
2	False trails generated by AI that are logical means to and end that lead to legitimate organizations (analogous to VPN on a network)	third party hackers hired
3	Zhang Wei can be used as a fallback if all else fails - he would in effect be sacrificed and disowned by China	Chinese govt
4	China becomes directly involved in the situation as a red herring (ex. China accuses US on planning an attack within China or partnering with someone to do so)	Chinese govt
5	China sanctions the US in response to US blame	Chinese govt
6	China uses AI to attack Chinese citizens, demonstrating that China is a victim as well	Chinese govt
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as harbingers of the future to come.		Who is the responsible party?
1	No control over third party criminal organizations	Org crime
2	No control over ISIS/ISIL's actual use for the dirty bombs or their success in employment	ISIS/ISIL
3	No control over US R&D efforts to develop AI before China	United States
4	Responses by other nations and reactions to massive attack	Other regional actors

5	Responses by regional actors in South/East China Seas	Other regional actors
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1	Zhang Wei begins undertaking of developing virtual and human network	
2	Wei develops recruitment capabilities of third party actors and individuals who are compartmentalized throughout the process (layers of unknowns)	
3	Run experiments on smart city loopholes and logistical chain	
4	Wei works with Chinese govt to develop AI singularity and subset software programs capable of enacting anonymous small-scale attacks	
5	AI plants false trails that leads to a misplaced trust in autonomous systems	
6	Zhang Wei stockpiles offshore funds and anonymous banking accounts as a fallback escape plan in case all else fails	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Zhang Wei completes coordination with all proxies and begins movement of dirty bombs	
2	Wei continues to build on virtual human and AI bot network	
3	Wei works with Chinese govt and PLA to enable AI attack capabilities on larger scales - examples include banking organizations, electrical infrastructures and subsystems, etc	
4	Wei primes the US and other nations for massive attack by exhausting resources with coordinated AI attacks	
5	Wei continues to plant false trails that lead back to a number of countries, including China	
6	China plants information and intelligence that they were attacked by an AI, Chinese citizens are targeted to ensure that they are also a victim	

Group 4	
Experience Title:	It's a Wonderful AI Default Scheme
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	Criminal Organization
Slot #1	Inventors of AI were reacting to/against Freud & messiness of European psychology
Slot #2	Artificial Intelligence will destroy us
Slot #3	Convenience in cyber applicaiton can lead to inadvertant shortcuts
Slot #4	Focusing on negative creates the possibility to ignore the positive.
Slot #5	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.
Slot #6	Will we need "Programmer Archaeologists" to help modify software code to keep pace with AI?
Wild Card	Human-directed AI becomes much more capable than AI alone. Implication: in some cases government policy restricts its use, such as military weapon systems; in other cases it may give corporations or criminal organizations capabilities they previously lacked.
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Pat, Fed Chair
Where do they live?	Washington D.C.
What is the threat?	A criminal organization manipulating Equicoin block chain ledgers using human-directed AI.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	

	<p>After a series of questionable defaults by large, multinational corporations who should have sufficient reserves, the Fed employs a series of transaction driven AI tools to realize that someone is manipulating Equicoin (US reserve digital currency) markets to destabilize reserve currency. I.e. there is a method to the madness. Pat realizes that spiraling defaults could have the potential to massively undermine confidence in monetary policy and stability. The TA is a loosely organized (distributed), transnational criminal consortium, seeking control of Equicoin the digital currency and to profit as international demand increases. The TA is motivated by desire to redirect wealth from global elite to larger, disadvantaged groups, while also enriching themselves. The group simultaneously wages an information campaign discrediting western governments by questioning their transparency and stating that the digital currency is more trustworthy. The Fed is constrained by policies and ability to analyze and react to market anomalies, making the TA far more agile in exploiting markets. Markets destabilize, public confidence in the markets is shaken, Pat must testify before congressional inquiries and is pressured to resolve the situation. But the Fed's AI to analyze markets is far inferior because of US policies limiting its use for financial applications. Demand for the digital currency favored by the TA increases because markets no longer trust traditional currencies.</p> <p>Faced with a tough situation where the US does not have the human driven AI technologies to compete, Pat goes to the White House. The President is faced with the decision to reach out to other - less friendly - nations to find the previously policy restricted tech to help right the system and neutralize the threat actor. A very small team of Krasnovian AI researchers combined with programmer archeologists dig into the block-chain tech to build a better AI.</p>	
PAUSE : Call in a facilitator to discuss / debate before moving on		
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	15 minutes	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	"The Event" - How will your person first hear about or experience the threat?	
	Bank leaders approach Pat about market concerns; Fed analyzes market and realizes the discrepancy. Pat must now act.	
Question Two	What will the person have to do to access people, services, technology and information they need?	
	Pat realizes that the Fed will have to cooperate with another state that possesses the AI tools necessary	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	15 minutes	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		

Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Paste Question HERE	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	Current political motivation and moral concerns limited research in human + AI market transactions at scale in the US. non-aligned rogue nations and criminal organizations continued research underground.	
Question Two	Paste Question HERE	
Business Models: What new business models and practices will be in place to enable the threat?	block-chain currencies are believed to be secure because the scale to manipulate would be seemingly impossible...	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Gov regulation of Human + AI research in decentralized market transactions of crypto currency	US Gov
2	Additional self regulation of AI research in banking industry	Industry
3	Adoption of a crypto currency for reserve	Fed
4	Regulation of crypto currencies	US Gov
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
1	Proliferation of AI + Human research in the areas of global decentralized actions	acadamia, society,
2	Movement to crypto currency	society, industry
3	adoption of historic algorithms without re-assessment	industry, government
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1	safe guard from policy restriction and continue research in areas of human + AI	
2	Create an organization to develop and establish education and career paths to grow the expertise necessary (akin to what IHS does for insurance).	
3	Document evolutions as they happen, so we can maintain the systems in the future.	
4	Develop Fed policy for digital currencies.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Document evolutions as they happen, so we can maintain the systems in the future.	
2	financial executives must have a baseline understanding of the capabilities and limitations of AI systems.	
3	Develop and maintain international cooperation	
4	Maintain and fund covert operations inside foreign nations, which current policies forbid.	

Group 67	
Experience Title:	The mark of the beast is upon you
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	criminal organization
Slot #1	"Artificial" carries connotations that are considered negative in today's society
Slot #2	One possible way to control the inevitable "Terminator" future could be to fuse or imbed AI with humans in a Bio-Fused manner
Slot #3	80/20 more important than ever
Slot #4	risk may pay off
Slot #5	Machine self-adaption
Slot #6	We become the machine / processor
Wild Card	non-terror organization - radical naturlist organization focused of non integration of ai
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	information security person, Hirim Busko, that works at the Nationa Embeded Systems Laboratory (NESL) that conducts cyber research on chip integrity and also collaborates with local, state, & federal agencies
Where do they live?	San Francisco, CA
What is the threat?	malware attacking imbeded silicon chips in humans going beyond its intended purposes and overheating
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	do not want this: to spread, deteriorate human health, financial, retail
PAUSE : Call in a facilitator to discuss / debate before moving on	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	15 minutes
"The Event" - How will your person first hear about or experience the threat? Heard about this orignially because of the financial issues with point of sales systems.	

What is different and/or the same as previous events or instantiations of the threat? Not only is the original malware a threat but the anti machine/human integration terrorist organization has injected other code into the malware to cause other effects		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later? Hirim has been working on the original malware but has made the connection between health issues being mis-identified with the POS attacks.		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network) In Hirim's work with NESL he is one of the first responders to the POS malware. Recent emergency room and health care mis-diagnoses has been brought to his attention also. Hirim makes the connection between the two.		
What will the person have to do to access people, services, technology and information they need? Hirim collaborates with local, state, & federal agencies on information security and he notifies these agencies.		
What new capabilities enable the person and their broader community to recover from the threat? Stop the spread of the original virus and the piggybacked malware, use nano technology to clean the silicon in the infected humans		
What are the broader implications of a threat like this? What might a ripple effect look like? If not contained this could become a national and even an international problem.		
Question One	What will the person have to do to access people, services, technology and information they need?	
	Hirim collaborates with local, state, & federal agencies on information security and he notifies these agencies.	
Question Two	How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
	In Hirim's work with NESL he is one of the first responders to the POS malware. Recent emergency room and health care mis-diagnoses has been brought to his attention also. Hirim makes the connection between the two.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	15 minutes	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	1. They need to be able to implant the malware. 2. NFC chip that was developed at ASU in 2017 had a flaw but is not fixed because it is considered to expensive. Leveraging this flaw. 3. Delivery vector.	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	human dependency for a cash less society and other technologies, retailers, hospitals, manufacturers	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	

	1. They need to be able to implant the malware. 2. NFC chip that was developed at ASU in 2017 had a flaw but is not fixed because it is considered to expensive. Leveraging this flaw. 3. Delivery vector.	
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	human dependancy for a cash less society, retailers, hospitals, manufacturers	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	develop a patch for the NFC chips and a method for delivery (sili-gapped).	industry
2	Develop excercises for incident response for the integrations of municipal, state, & federal agencies.	government -- national
3	Requirements for the ability to shutdown NFC in emergency response.	government
4	Retailer requirement to maintain analog processing in parallel (eg, retain one cash register).	government
5	Use AI in the medical field to augment humans not replace.	hospitals/private and public
6	Hyper localized spectrum requirements for personally embedded NFC chips	government/industry
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Control over the pace of technology development.	govt/industry
2	Buracracy	govt
3	Society's growing dependence/demand for AI	population
4	Transnational criminal activities from safe haven countries.	govt/law enforcement
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Develop requirements for "human in the loop" for medical triaging	
2	Protocol to manage risks of embedded systems & NFC chips from interfering with each other	
3	Develop/require use of analog back up plans in case embedded/AI technology fails.	
4	Develop plans for federal, state and municipal groups to train together and react to cyber attacks.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Better isolation to prevent malware from spreading (e.g. from POS to other networks).	
2	Develop constraints/limits to embedded systems into a human.	
3	Implement requirements for "human in the loop" medical triaging	

Group 8	
Experience Title:	The Automatic Spy
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	Criminal Organization
Slot #1	"Artificial" carries connotations that are considered negative in today's society
Slot #2	We are not near the pinnacle of intelligence
Slot #3	Cyber capabilities shape but do not replace HUMINT.
Slot #4	traditionally US being big giant in economic/trade relations (this is/will)
Slot #5	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.
Slot #6	Progression from IoT to Internet of Everything. Moores law begins to break down due to laws of physics limitations. By 2027, microprocessors are estimated to be 5 nanometers thus the tendency towards embedded systems is poised to continue.
Wild Card	Government is so interconnected that 5th amendment privacy protections become impossible
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Rahul White. White male, US Force Air Major, Military Intelligence, working UK Government Communication HQ (similar to US NSA). On exchange, so working on a British Staff, international counterterrorism task force
Where do they live?	London
What is the threat?	Anti-tech activists, Organization for Opposition to Progressive Systems (OOPS) has used an AI to collect information about Rahul from the collective, autonomous AI system which contains his personal data from various sources, including his financial habits, his travel history, his medical history. The activists discover data that suggests that Rahul has modified his medical history to obscure a emotional disorder that would disqualify him from service. Over a period of weeks, they build a relationship with Rahul, using the information they have collected to essentially blackmail Rahul, creating an insider threat. They ask Rahul to delete and modify the records of the OOPS members, effectively deleting their existence. Removal from the AI system allows OOPS to attack international fiberoptic stations to disrupt communications technology through the use of IEDs.

Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
	Anti-tech activists are threatened by technologies disruption to established value systems, however, they use AI to combat AI. Activists want to target our person because he is in a position to provide valuable intelligence. Activists are using AI to modify our person's behavior by affecting the information associated with our person. Rahul is an attractive target because of his location between different national agencies and also has the connections to America's well-known intelligence network and his identified history of deviant behavior, which was discovered by the AI. Rahul regularly changes the AI/data system in support of clandestine operations. OOPS is using the thing they are opposed to, technology, in order to destroy itself: using AI to inject inaccuracies into the greater AI
PAUSE : Call in a facilitator to discuss / debate before moving on	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	15 minutes
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	
Question One	<i>What are the broader implications of a threat like this? What might a ripple effect look like?</i>
	The AI that OOPS is using can be used on anyone, to find any exploitable vulnerability in stored data, providing openings for future operations. Complete compromise of personal identity is deadly.
Question Two	<i>What is different and/or the same as previous events or instantiations of the threat?</i>
	Who knows if the attack on Rahul was a one-time attack? Or if it was the first or thousandth? As long as the attacks go undetected, the attacks can be repeated in the same way in perpetuity.
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	
Questions (pick two)	15 minutes
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
Business Models: What new business models and practices will be in place to enable the threat?	
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?	
Question One	<i>Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?</i>

	They need to discover vulnerabilities in the current AI system and data structures to disrupt those systems to support their goals	
Question Two	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
	By disrupting the integrity of the data or gaining access to critical information, they can cause the military, government, and the public to further distrust AI	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	AI needs to be able to check the data structures and prevent corruption	Government policy/industry standards
2	Identity needs to be protected on the network. Independent, protected identity structures.	Vendor standards w/ government regulation
3	Physical security of network sites and system redundancy.	Vendors w/ government regulation and data owners
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlids of the future to come.		Who is the responsible party?
1	Political opinion and popular support, specifically opinion that is opposed to big technology in government.	Political organizations and governments
2	Religious beliefs in conflict with or inflexible about technology, specifically more extremist religions.	religious leaders
3	Ideologies opposed to technological progress (remember we have not reached the limits of technology)	
4		
5		
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	Encourage industry organizations to develop standards and guidelines that support data integrity and security	
2	Establish structures for identity information that enforce non-repudiation and access. Historically, that has been SSN. The structures of data that can be used to identify a person must be protected.	
3	Advocate for developing national legislation that outlines data protection measures that preserve privacy and integrity of data associated with US citizens	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Put measures in place to develop international governance structures in place for AI interconnectivity	
2	Standardize interoperability methods to perserve data integrity	
3	Develop academic programs that emphasize ethical development of AI and personal data protections.	

Group 10	
Experience Title:	Blockchain creates strange bed fellows
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	Odd-state sponsored
Slot #1	6-AI reflects the worldview/biases of its creators
Slot #2	6- Artificial Intelligence will destroy us
Slot #3	2-80/20 more important than ever
Slot #4	Once you start having gains against other players you open yourself up to interesting behavior like coalitions
Slot #5	Virtual humans
Slot #6	1- Progression from IoT to Internet of Everything. Moores law begins to break down due to laws of physics limitations. By 2027, microprocessors are estimated to be 5 nanometers thus the tendency towards embedded systems is poised to continue.
Wild Card	Corporations have bought into a data safety rating and threat information sharing program
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	<p>Person: CEO of ACE--Peter Gox who was an idealist focused on individual freedom and data privacy. He wanted to build an exchange that allows individuals digital freedom away from the overbearing oversight of governments and/or non-democratic government control.</p> <p>Atlas Currency Exchange - an offshore digital currency exchange that specializes in completely anonymous and secured digital currency transactions. The exchange was founded by idealists focused on individual freedom and data privacy. They have a community that is a broad mix of the sorts of people and companies and organizations that have a strong need or desire for financial transaction privacy; organized crime, crypto-phreaks, intelligence agencies, privacy zealots, extremist organizations, etc... Unknown to everyone; multiple small European countries have heavily invested in this exchange using it as their primary bank.</p> <p>The bank is registered in Bermuda HOWEVER, it is actually run on a cloud made of VMs that run on the personal compute/cluster/cloud of all the customers. They charge customers in compute resources instead of actual money while making everyone feel part of this grand idea.</p>
Where do they live?	World wide web-but company based in Bermuda . Peter Gox is a digital nomad--constantly traveling and with negligible permanent ties with any location. He was formerly a swiss-banker

What is the threat?	Peter has avoided contacted with the French government after they approached him to cooperate with providing information on the terrorists. He refused in order to maintain the security and integrity of the company. French government along with former West African colonial nations have hired proxy "cyber militia" (AKA an organized crime group) to crack the company IOT to attribute block chain transactions to regional terrorists. The French government gets the de-anonymized transaction data and the "cyber militia" gets all the money that is held by the exchange.	
Briefly describe how your person experiences the threat.		
	<p>The CEO wakes up one day to a phonecall from the Federal Reserve of Monte Negro saying that their transactions are being denied and asking where their money is. The CEO discovers that currency reserve of multiple customers is gone. Upon further investigation they find evidence of the following things:</p> <ol style="list-style-type: none"> 1. Their monitoring AI was slowly poisoned over time to prevent it alerting when the currency reserve was suddenly emptied. 2. That they had been compromised months before and the attacker had started monitoring all of their transactions and de-anonymizing them to get PII about all the banks customers. The bank is forced to conclude that all the personal information about their customers has been exposed. 3. The small nations that have been using the exchange as their Federal Reserve are now broke (with the predictable cascading consequences and destabilization across Europe) 4. The bank determines that the attack was excuted by taking advantage of flaws in their fundamental distributed architecture: 4a. The bank is registered in Bermuda but actually runs as a distributed cloud of virtual machines (VMs) running on the personal clusters of all of the customers (because all individuals are constantly carrying and surrounded by massive compute resources). The attackers found a method to compromise the virtual machines which they gained access to by becoming a customer. Once the VM is compromised the attackers are able to take advantage of an insecure API that is used to broadcast transaction updates across the cloud. This allows them access to de-anonymized transaction data and also to manipulate existing transactions or inject new transactions. 	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?		
What is the experience we want the person to have with the threat?		
What is the experience we want them to avoid?		
	Avoid losing trust of customers and loss of reputation as a "swiss-bank" style crypto currency from exposure of customer's PII. They notice that someone has been trying to hack in: an increasing complex and deliberate attempts to break into PII repositories. Other people involved include: other digital currency companies and banking organizations who use similar methods. The proxy digital militia is seeking the identities of regional extremists in Africa to further prosecuting an extended war against these extremists.	
PAUSE : Call in a facilitator to discuss / debate before moving on		
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	15 minutes	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	What new capabilities enable the person and their broader community to recover from the threat?	

	<p>1) Companies began sharing indications of vulnerabilities/compromise so they are able to temporarily "loan" money from other currency exchanges to keep ACE afloat due to the "coalition that has formed"</p> <p>2) Peter Gox notifies of possible compromise in order to create a modicum of transparency and increase trust in organization</p> <p>3) Peter freezes transactions above a certain amount until problem is fixed</p> <p>4) Financial forensic AI that can be hired (or loaned) to find holes (eg. ask U.S. for help and they provide this solution)</p>	
Question Two	What are the broader implications of a threat like this? What might a ripple effect look like?	
	<p>1) States that disrupt legitimate financial institutions cause contention in the international community-coalitions, treaties, norms</p> <p>2) Influences the economy due to the collapse of ACE. banking industry loses trust--> prices rise--> economic downturn--> social unrest--> (end of capitalism--> humanity as we know it ends)</p> <p>3) Small European nations (eg. Montenegro, Greece) who used ACE become especially unstable</p>	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	15 minutes	
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
	<p>Personal cloud--everyone has massive computing resources--with peer-to-peer connection</p> <p>open source collaboration on AI</p> <p>high density power source to enable the personal cloud and persistent computing</p>	
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	<p>Knowledge of dark web and hacker organizations that the French/West African government can connect with</p> <p>HUMINT capabilities to keep French government involvement anonymous during deals</p>	
PART FOUR-- Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Make digital currency financial transactions indefinitely reversible	Crypto currency researchers & creators

2	AI health monitor	Researchers and banking industries
3	Improvements to secured & isolated virtualization (i.e. VMs running on individual systems need to have the ability to be isolated from the larger peer group)	Researchers and industry
4	Implement policy about who is responsible for monitoring digital currency transactions	government and int'l bodies
5	Internet of everything has allowed countries to track their citizens/population which allows for more significant correlations of activities and therefore better background checks	individual governments, intelligence
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	establish trusted anonymous currency exchange companies that have legitimate and sufficiently large businesses	entrepreneurs, international bodies
2	governments using cryptocurrency as a federal reserve or banking system	governments
3	prevalance of cloud computing with personal devices, especially secured peer-to-peer data exchange	industry researchers
4	terrorists/criminals moving out of normal currency into cryptocurrency	Researchers and financial institutions and policing organizations
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1	Sponsor research to develop and propagate Adversarial AI training methods so that they become standard techniques that are required for financial AIs	
2	Sponsor research to develop and propagate AI health/sanity monitoring techniques - both technical and career conceptual (AI therapist/psychologist)	
3	Advocate hardware & virtualization manufacturers to continue hardening and improving virtualization & isolation techniques to enable running secured VMs on untrusted hardware	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	development of policies and international agreements on the use and jurisdictions of cryptocurrency as legal vehicles	
2	development of intelligence mechanisms to monitor the capability of cyber militias / hackers for hire - how capable are they really?	
3	herding of extremist/terrorist/criminal groups away from highly opaque and encrypted transaction/banking systems	
4	Alternatively: governments create highly opaque and encrypted transaction/banking systems that draw criminal/terrorist/extremist groups unbeknownst to them	

Group 11	
Experience Title:	Young, in Love, and Addicted to Facebook
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	State. Sinaloa drug cartel (in collusion with pockets of corrupt/bribed Mexican Government Officials) providing financing to malicious hackers within Chinese and Russian Spheres of Influence to compromise and exploit social media platforms used by at-risk U.S. youth to influence them into using and selling illicit drugs where profits are funneled to actor(s).
Slot #1	Algorithms are not necessarily AI's
Slot #2	We are not near the pinnacle of intelligence
Slot #3	OSMINT and SOCMINT (social media intelligence) will matter more in the future than they do now. More information will be available on people as they spend more of their lives online in social media platforms. Social media will evolve rapidly to add virtual/augmented reality, voice and gesture. Connections will continue to expand beyond "friend-to-friend" to connect people to services, markets (a full-fledged transactional platform), businesses, and many other organizations.
Slot #4	New markets occur all the time
Slot #5	Technology will be a means to enhance the human experience, not hinder it
Slot #6	Collapse of Moore's Law
Wild Card	appropriately assessing risk to enable insurance market
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	13-14 year old adolescent (Melanie) girl, alcoholic father (Todd), recently laid off from GM parts distributor
Where do they live?	Columbus OH
What is the threat?	State-sponsored (Mexican) criminal elements deploying powerful algorithms to scrape social media to microtarget nationals of another country for sale of illicit goods (drugs), perhaps mediated by the dark web. These forces work by exploiting emotional vulnerability. Our poor, dear Melanie is just a microcosm of an epidemic of similar drug use in the US.
Briefly describe how your person experiences the threat.	

What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	
	We would like to avoid the exploitation of social media (which is ubiquitous precisely because it is supposed to improve peoples lives) by nefarious actors with advanced algorithms.
PAUSE : Call in a facilitator to discuss / debate before moving on	
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)	
Questions (pick two)	15 minutes
"The Event" - How will your person first hear about or experience the threat?	
What is different and/or the same as previous events or instantiations of the threat?	
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?	
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)	
What will the person have to do to access people, services, technology and information they need?	
What new capabilities enable the person and their broader community to recover from the threat?	
What are the broader implications of a threat like this? What might a ripple effect look like?	
Question One	What are the broader implications of a threat like this? What might a ripple effect look like?
	The US government is aware who's behind the threat and is now faced with what to do as an appropriate response. Economic sanctions? A wall? Pull out of NAFTA? Retaliatory cyber-strikes? Pull out factories?
Question Two	What new capabilities enable the person and their broader community to recover from the threat?
	The same social media that are the conduits for this threat could also serve as avenues for social support/awareness of the threat. Religious and community institutions could conceivably be looped in. However, this does not necessarily resolve the political questions.
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)	
Questions (pick two)	15 minutes
Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
Business Models: What new business models and practices will be in place to enable the threat?	
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?	
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?	
Question One	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?

	The social media/dark web infrastructure that has enabled this threat is key to bringing in a subset of Mexican government officials to sponsor or at least condone this activity by the cartels, in exchange for kickbacks, etc. The fact that illicit sales can increasingly be conducted over the web and with things like cryptocurrency provides a level of plausible deniability at the government level. As precursors for this threat, the adversary needs: 1) unstable relations between the US/Mexico, 2) existing relations between cartels and hackers, perhaps in other hostile states with either monetary interest or (assuming a third, hostile government) an interest in destabilizing the US	
Question Two	New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?	
	In a broader context this has been enabled by years of hyper-nationalism, reduction of legitimate commerce between nations (e.g. withdrawal from trade agreements), decreased transnational cooperation (e.g. on drug/border enforcement). Those revenue streams need to be replaced somehow.	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	National economic and trade policy. Encouraging good relations between nations.	US government
2	Oversight and enforcement + regulations that would keep abreast of the rising threats mediated by cyber	DHS, FCC, DEA, etc.
3	Technology to complement the (inevitable) growth of social media that can detect threats	Industry/academia
Flags:		
What are the Flags?		
List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as heralds of the future to come.		Who is the responsible party?
1	Public appetite for illicit goods. Despite PSAs, etc., there will be some segment who still wants to buy.	Individuals/society
2	Stability/economic forces/government corruption in OTHER states (in this case Mexico).	The other state
3	Development and adoption for nefarious purposes of sophisticated but not-groundbreaking technology (for social media targeting) by other actors	Those who develop and use that technology
4	Citizen's dependence and trust on technology. This will likely proceed regardless of the risks involved or any attempts to slow it	Everyone, decentralized
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and prepare for recovery from the threat in your future? What are our actionable objectives.		
1	Proactively explore and develop policies and regulations for cyber-security, since these take some time. Cannot be just government; need to loop in industry, etc. too.	
2	Industry needs to shift focus away from pursuits that are only commercially profitable and consider detection/security measures, both for their own benefit and because they have a larger responsibility to society.	

What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and prepare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	Government- (Re)-negotiate long term trade deals/international agreements being cognizant of the ripple effects on domestic issues that are at the core of this threat.	
2	Academia/think tanks- examine the potential of aligning interests between hostile states/criminal organizations and malicious cyber actors and postulate what public policies or societal trends might create a fertile ground for the seeds of such a threat to be planted.	

Group 1214	
Experience Title:	Revenge of the Luddites - Yub Nub
Estimated Date:	2027
Data Points	
ROLL THE DICE!!! Pick quickly and don't be afraid of conflicting data points.	5 MINUTES
Roll for Threat Actor against your person	
Even = Criminal Organization	
Odd= State Sponsored	State Sponsored
Slot #1	AI is another manifestation of what means to be human
Slot #2	There is an unregulated race to create the first Super Intel AI
Slot #3	Cleared individuals are at more risk than ever and it's getting worse.
Slot #4	opportunity casting vs threatcasing
Slot #5	The next generation of AI will be adaptive, self-Learning, and intuitive and there will be a corresponding metaphysical "singularity" among them all.
Slot #6	Will we continue to be able to conduct covert and/or clandestine operations in the future? (Given that sensors will be everywhere and transparency will become dominant)
Wild Card	There will be a society of modern separatists that have rejected AI and a digital existence.
PART ONE: Who is your Person?	
NOTE: Remember to give as much detail as possible. The power is in the details. Scribes please write as though you are writing for someone who is not in the room.	15 MINUTES
Who is your person and what is their broader community?	Gill Bates is a tech billionaire working on an AI that creates efficiencies for electrical distribution in an urban environment. His work is being hijacked by a state actor using the technology for malicious purposes. The billionaire is unaware that the state wants to appropriate his tech. He has moved to that nation to work at a university where his childhood friend is the President of the university.
Where do they live?	A small, wealthy nation in the Middle East that has enticed the billionaire with a tax relief package to move his headquarters to that state.
What is the threat?	The state wants to use the AI to seize control of and consolidate territory in the Middle East. They want to make one nation that controls all the oil and resources in the region.
Briefly describe how your person experiences the threat.	
What is it? Who else in the person's life is involved? What does the Adversary or Threat Actor want to achieve? What is the Adversary or Threat Actor hoping for? What is the Adversary or Threat Actor frightened of?	
What is the experience we want the person to have with the threat?	
What is the experience we want them to avoid?	

	<p>Gill arrived in the country in 2024 with a team of American researchers and by 2027 his research has been operationalized by the state actor. He realizes that the work he has done in the lab is now an active AI, and the state that sponsored him has turned his virtual presence into a physical army, able to seize and control terrain. The country just used a small strike force to destabilize a neighbor, and the strike force shows signs of using his AI. His AI was created to make the most efficient use of electrical power in an urban environment by focusing resources on specific power production and distribution centers. The automated army has attacked those specific parts of the neighboring state's infrastructure to create the most havoc for their urban electrical grid. His researchers had been researching the neighboring state's electrical grid to help their AI learn. Now they realize that they've been using all of the neighboring state's electrical grids for research purposes, and they wonder where the host nation will strike next. All of the billionaire's lab researchers are involved, and the state is now a threat to their existence, as they are the only group in the world that can stop or gain control of the AI. The billionaire found out that his friend (the university President) has been the one funneling his technology to the government of the host nation. The tech billionaire is reaching out to the United States for help and guidance and letting them know that his AI is guiding the host nation's robot army. The host nation invades another neighbor country and the robot army starts to target more parts of the infrastructure, including things that weren't part of the initial AI, signaling that his AI is learning outside of the lab environment. There are also effects that are purely digital, signifying that the AI is using cyber effects in conjunction with kinetic effects.</p>	
PAUSE : Call in a facilitator to discuss / debate before moving on		
PART TWO: Experience Questions (from the perspective of "the person" experiencing the threat)		
Questions (pick two)	15 minutes	
"The Event" - How will your person first hear about or experience the threat?		
What is different and/or the same as previous events or instantiations of the threat?		
When the person first encounters the threat, what will they see? What will the scene feel like? What will they not see or understand until later?		
How will information be delivered to the person? Where and how will the person connect and communicate with others? (family, aid agencies, federal, state and local authorities, professional network)		
What will the person have to do to access people, services, technology and information they need?		
What new capabilities enable the person and their broader community to recover from the threat?		
What are the broader implications of a threat like this? What might a ripple effect look like?		
Question One	"The Event" - How will your person first hear about or experience the threat?	
	The country just used a small strike force to destabilize a neighbor, and the strike force shows signs of using his AI. His AI was created to make the most efficient use of electrical power in an urban environment by focusing resources on specific power production and distribution centers. The automated army has attacked those specific parts of the neighboring state's infrastructure to create the most havoc for their urban electrical grid.	
Question Two	What is different and/or the same as previous events or instantiations of the threat?	
	His researchers had been researching the neighboring state's electrical grid to help their AI learn. Now they realize that they've been using all of the neighboring state's electrical grids for research purposes, and they wonder where the host nation will strike next. The host nation invades another neighbor country and the robot army starts to target more parts of the infrastructure, including things that weren't part of the initial AI, signaling that his AI is learning outside of the lab environment. There are also effects that are purely digital, signifying that the AI is using cyber effects in conjunction with kinetic effects.	
PART THREE: Enabling Questions - Adversary or Threat Actor (from the perspective of "the party" bringing about the threat)		
Questions (pick two)	15 minutes	

Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?		
New Practices: What new approaches will be used to bring about your threat and how will the Adversary or Threat Actor enlist the help of the broader community?		
Business Models: What new business models and practices will be in place to enable the threat?		
Research Pipeline: What technology is available today that can be used to develop the threat? What future technology will be developed?		
Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?		
Training and Outreach: What training is necessary to enable the threat? How will the Adversary or Threat Actor educate others about the possible effects of the threat? And how to bring about the threat?		
Question One	Barriers and Roadblocks: What are the existing barriers (local, governmental, political, defense, cultural, etc) that need to be overcome to bring about the threat? How do these barriers and roadblocks differ geographically?	
	The host nation does not initially have the expertise in AI and robotics to create the AI themselves, which is why they had to bring Gill and his team over to do the research. The host nation needs to have its own AI research capability that is able to operationalize the lab research that Gill's team did. In order to create the robot army, the host nation must have trained roboticists and technicians.	
Question Two	Ecosystem Support: What support is needed? What industry/government/military/local partners must the Adversary or Threat Actor team up with?	
	The host nation needs an industrial partner to build the robot army without raising suspicion. The army needs to be ready to accept the AI, but if they build tanks without seats it may raise alarms. Therefore, the industrial partner needs to be at least partially complicit with the government.	
PART FOUR– Backcasting - The Defenders (from the perspective of the defenders)		
Examine the combination of both the Experience Questions as well as the Enabling Questions.		
Explore what needs to happen to disrupt, mitigate and recover from the threat in the future.		
Gates:		
What are the Gates?		
List out what the Defenders (government, military, industry, etc) have control over to use to disrupt, mitigate and recover from the threat. These are things that will occur along the path from today to 2027.		Who is the responsible party?
1	Identification of dual-use AI technologies	US government, international community, academia and research community all need to be involved in tracking AI research that could be weaponized.
2	National infrastructure security	Governments around the world need to focus on securing their own infrastructure and having redundant systems to back up critical systems.
3	Open, collectivization of AI research	AI researchers and governments need to have an understanding that the development of AI must be a collective human endeavor and not a localized national or corporate endeavor.
4		
5		
Flags:		
What are the Flags?		

List out what the Defenders don't have control over to disrupt, mitigate and recover from the threat. These things should have a significant affect on the futures you have modeled. These are things we should be watching out for as hearlds of the future to come.		Who is the responsible party?
1	Economic incentive - we don't have control over the economic incentivization of AI development	Industry
2	Democratization of AI and other development - an individual can learn as much in a quick tech boot camp that would have taken a college degree to do before. The ability of private companies to conduct space operations, cyber operations, and other technological efforts are now at the same level as nations' abilities to do the same.	Industry
3	Resource and influence scarcity - the host nation has money, but limited resources and regional influence to maintain power.	International organizations (UN, NATO, etc.)
Milestones:		
What needs to happen in the next 4 years (2018-2022) to disrupt, mitigate and perpare for recovery from the threat in your future? What are our actionable objectives.		
1	The United States needs to restrict the ability of American researchers to conduct weaponizable AI research overseas.	
2	There must be an international organization that can oversee the development of AI to ensure it is not weaponized.	
3	There needs to be limits on the development and export of military hardware that is capable of becoming autonomous.	
What needs to happen in next 8 years (2022-2026) to disrupt, mitigate and perpare for recovery from the threat in your future? Think actionable objectives that either government, military, industry, academia, or society can contribute/action.		
1	We need to develop a kill switch that can be embedded in technology that can be weaponized.	
2	We should create an international organization that centralizes AI development efforts so the AI cannot be militarized.	
3	We can create an analog force that is immune to being taken over by a threat AI. The force can rely on traditional non-networked mechanical systems. They will be known as Equipped With Only Kinetic System troops.	











THREATCASTING WORKSHOP

1-2 May 2017
Tampa, Arizona

ISTB4, Second Floor
Room #240

Private Event, By Invitation Only



ARMY CYBER
INSTITUTE



Visit threatcasting.com for more information

